Distributed Load Balancing with Nonconvex Constraints: A Randomized Algorithm with Application to Electric Vehicle Charging Scheduling

Lingwen Gan, Ufuk Topcu, and Steven Low

arXiv:1401.7604v2 [math.OC] 11 Apr 2014

Abstract

With substantial potential to reduce green house gas emission and reliance on fossil fuel, electric vehicles (EVs) have lead to a booming industry, whose growth is expected to continue for the next few decades. However, EVs present themselves as large loads to the power grid. If not coordinated wisely, the charging of EVs will overload power distribution circuits and dramatically increase power supply cost. To address this challenge, significant amount of effort has been devoted in the literature to schedule the charging of EVs in a power grid friendly way. Nonetheless, the majority of the literature assumes that EVs can be charged intermittently at any power level below certain rating, while in practice, it is preferable to charge an EV consecutively at a pre-determined power to prolong the battery lifespan. This practical EV charging constraint is nonconvex and complicates scheduling.

To schedule a large number of EVs with the presence of practical nonconvex charging constraints, a distributed and randomized algorithm is proposed in this paper. The algorithm assumes the availability of a coordinator which can communicate with all EVs. In each iteration of the algorithm, the coordinator receives tentative charging profiles from the EVs and computes a broadcast control signal. After receiving this broadcast control signal, each EV generates a probability distribution over its admissible charging profiles, and samples from the distribution to update its tentative charging profile.

We prove that the algorithm converges almost surely to a charging profile in finite iterations. The final charging profile (that the algorithm converges to) is random, i.e., it depends on the realization. We characterize the final charging profile—a charging profile can be a realization of the final charging profile if and only if it is a Nash equilibrium of the game in which each EV seeks to minimize the inner product of its own charging profile and the aggregate electricity demand. Furthermore, we provide a uniform suboptimality upper bound, that scales $O(1/n)$ in the number n of EVs, for all realizations of the final charging profile.

I. Introduction

Electric vehicles (EVs) are propelled by electric motors powered by rechargeable battery packs, while conventional vehicles run on internal combustion engines. EVs enjoy a higher energy efficiency, more environmental friendliness, a superior performance (in terms of noise, acceleration, maintenance, etc.), and less energy dependency on foreign fossil fuels than conventional vehicles [1]. Therefore EVs are promoted by the US government with tax incentives—a $7500 federal income tax credit for every EV (small neighborhood EVs excluded) purchased in or after 2010 [2]. Besides government support, EVs are much more fuel economic, e.g., the 2014 Honda Fit EV is estimated to cost $0.87 electricity per 25 miles drive [3], while the 2014 Honda Fit is estimated to cost $3.25 gas per 25 miles drive.[1]

[1]Assume that a 2014 Honda Fit burns 1 gallon gas per 30 miles drive, and each gas costs $3.90.

Such government support and fuel economy have lead to a booming EV industry: over 170000 highway-capable EVs have been sold in the US from 2008 to 2013, and 16 EV models rom 9 major manufacturers are available in the US market by March 2014 [4]. This trend is forecast to speed up as major vehicle manufacturers consecutively announce their EV plans [5]–[7].

A multitude of challenges, not only economical but also technical, have to be overcome to welcome the new era of EVs. On the economical side, EVs are currently sold at a much higher price than conventional vehicles, mainly due to the expense of battery packs. This is starting to change since battery prices are dropping by 20%–30% each year [8], and therefore EVs are likely to have a competitive price in the near future. Technical challenges are more fundamental: among other things, EVs have a limited range due to small battery capacities, and it takes hours to get the batteries fully charged. Except for Tesla Model S, all other currently available EV models have a range of at most slightly exceeding 100 miles [3]. Typically, it takes 3.6–12 hours to fully charge a battery, depending on the battery capacity and the charging technology [3].

The technical challenge that we seek to solve in this paper is the integration of EVs to the power grid. While EVs present themselves as large loads to the power delivery circuits, the power delivery circuits are designed something like five decades ago without bearing in mind the prosperous of EVs. Furthermore, EVs are likely to be charged after their owners arrive home from work since large-scale public charging facilities are still to be constructed. If not coordinated wisely, the charging of EVs may lead to coincidence peaks in electricity demand [9]. Consequently, power transmission/distribution lines carry much larger currents and power transformers are loaded much more heavily. As a result, circuit device lifespan will be greatly reduced [10], and voltages throughout the network may suffer from severe deviations from nominal values [11]—note that electricity appliances like air conditioners and televisions only work under well-regulated voltages.

On the other hand, if the charging of EVs is well coordinated, not only will some of the integration challenges be mitigated, but also ancillary services can be provided to stabilize the voltages and frequency [12]–[14], and power supply cost can be reduced [15]. For example, if the charging of EVs is coordinated such that the aggregate electricity demand is flat, then high peak electricity demand can be avoided. Consequently, circuit devices will be least loaded and their lifespans will be enlarged; voltages throughout the network will be less fluctuating and easier to regulate; electricity demand is easier to follow since it ramps up and down much slower, and therefore expensive power plants for following the demand can be reduced, bringing down the power supply cost. As another example, the charging of EVs can be coordinated to respond to network contingencies such as the unexpected shutting down of a power plant. In such cases, the charging of some EVs can be postponed to rebalance the electricity supply and demand, which is essential for preventing electricity blackouts.

Literature review

A variety of algorithms have been proposed in the literature to coordinate the charging of EVs, ranging from centralized algorithms where a coordinator makes decisions for the EVs [15]–[18] to distributed algorithms where each EV makes their own decisions [19]–[22]. Note that distributed algorithms may still require coordinators to

facilitate the communication among EVs.

Centralized algorithms are mainly for cost-benefit analysis purposes when the number of EVs is large, since the computation burden would be too heavy for a single computation unit. In [15], a large number of operational distribution networks in The Netherlands are investigated to quantify the impact of EV charging on various network levels. Results show that controlled charging can reduce infrastructure update investment by half over uncontrolled charging. In [16], uncontrolled charging and smart charging of EVs are compared empirically to highlight that a second peak electricity demand during night can be avoided by adopting smart charging. In [17] and [18], centralized linear programmings are proposed to compute the optimal charging profiles of EVs. In [17], the linear programming aims to minimize the power supply cost subject to circuit capacity constraints and vehicle owner's requirements. The linear programming in [18] further includes a power network physical model that captures voltage deviations.

By distributing the computation burden among different units, a distributed algorithm is more suitable for scenarios where a large number of EVs need to be coordinated. In [19], a distributed algorithm is proposed to schedule the charging of EVs such that the aggregate electricity demand is made flat. It is proved in [19] that when the EVs are identical, the obtained aggregate electricity demand will be optimal in the sense that it is as flat as possible. This notion of optimal aggregate electricity demand is formalized in [20]. Besides, a different distributed algorithm is proposed in [20] to schedule the charging of EVs. Furthermore, the algorithm proposed in [20] always obtains optimal aggregate electricity demand. The algorithms proposed in [21], [22] take another perspective: instead of trying to flatten the aggregate electricity demand, they seek to minimize the charging cost given a pre-determined electricity price profile by solving a dynamic programming at each EV.

All aforementioned works make the assumption that the battery of an EV can be charged intermittently at any power level below certain rating, e.g., if Single-Phase Level II charging is deployed, then the rating is 3.3kW [23] and an EV can be charged at any power level from 0 to 3.3kW when it is plugged in. However in practice, to prolong the lifespan of batteries, it is preferably to charge the batteries consecutively at a fixed pattern until they get fully charged [24]. Then, the only flexibility in scheduling the charging of EVs is to determine when the charging of each EV starts. Such flexibility can no longer be described by linear constraints. Moreover, one can show that such flexibility imposes nonconvex constraints on the scheduling of EV charging, calling for a different approach.

Summary of contributions

This paper proposes a distributed algorithm to schedule the charging of EVs in order to shape the aggregate electricity demand (e.g., to flatten the aggregate electricity demand or to track some target demand profile), in the presence of convex or nonconvex charging constraints (recall that charging constraints are nonconvex if the only flexibility in scheduling is the time when an EV starts charging). In particular, our contributions are threefold.

First, in the special case where charging constraints are convex, we improve the convergence rate of the distributed algorithm proposed in [20] by updating the charging profiles of EVs at different speeds based on their total energy consumption. A load balancing problem that aims to minimize the ℓ_2 norm of the aggregate load profile is formulated in Section II-A, and the optimal EV charging scheduling problem formulated in Section II-B is an example of the

load balancing problem. In the special case where charging constraints are convex, a distributed algorithm is presented in Section III-A to compute the optimal charging schedules of EVs.

This algorithm, referred to as Algorithm 1, is similar to the gradient projection algorithm (GPA) proposed in [20]. Both algorithms are iterative, assume the availability of a coordinator that communicates with all EVs, and share the same communication pattern. In each iteration of Algorithm GPA, the coordinator receives tentative charging profiles from EVs and computes the partial derivatives of the objective function with respect to each tentative charging profile. The partial derivatives are identical since the objective function only depends on the aggregate of charging profiles, and is therefore symmetric (the interchange of two arbitrary charging profiles does not change the objective value). Then, the coordinator broadcasts the common partial derivative to all EVs. After receiving the common partial derivative, each EV updates its tentative charging profile by first taking a step along the negative direction of the partial derivative and then projecting the resulting charging profile back to the set of admissible charging profiles. It has been proved in [20] that the sequence of tentative charging profiles obtained by Algorithm GPA always converges to the set of optimal charging profiles.

Algorithm 1 improves over Algorithm GPA in that the step sizes, taken by the EVs to update their tentative charging profiles, are chosen to speed up the convergence rate. The intuition of choosing the step sizes is that, EVs with larger total energy requests should update their charging profiles faster since they have a larger set of admissible charging profiles to optimize over. In particular, the step sizes are chosen in proportional to the total energy requests of EVs.

Second, in the general case where charging constraints are nonconvex but certain technical conditions hold, we propose a randomized distributed algorithm for scheduling the charging of EVs. This algorithm, referred to as Algorithm 2, defers from Algorithm 1 only in how EVs update their tentative charging profiles. Recall that in Algorithm 1, EVs update their tentative charging profiles by first taking a step in the direction of decreasing the objective value, and then projecting the resulting charging profile back to the set of admissible charging profiles. However when charging constraints are nonconvex, the set of admissible charging profiles can be difficult to project onto. Moreover, one can prove that any deterministic algorithm is guaranteed to diverge in the simplest setup where EVs are identical, unless the communication pattern is different from Algorithm 1. This highlights the necessity of a randomized algorithm to handle the nonconvex charging constraints.

In Algorithm 2, to update its tentative charging profile, an EV first computes a charging profile $\hat{x}$ in the convex hull of admissible charging profiles, then finds a probability distribution over admissible charging profiles that has expectation $\hat{x}$, and finally samples from the probability distribution to update its tentative charging profile. Noteworthy, Algorithms 1 and 2 share the same coordinator operations: the coordinator receives tentative charging profiles and computes a broadcast signal—the common partial derivative of the objective function with respect to individual charging profiles. Such compatibility between Algorithms 1 and 2 allows for the joint scheduling of EVs with convex and nonconvex constraints: for EVs with convex constraints, run the EV algorithm in Algorithm 1; for EVs with nonconvex constraints, run the EV algorithm in Algorithm 2.

The derivation of Algorithm 2 is based on the martingale theory [25]. Loosely speaking, a supermartingale is a

stochastic process that decreases in expectation, and all lower bounded supermartingales converge almost surely. Algorithm 2 is designed such that the objective values obtained in consecutive iterations form a supermartingale. More specifically, the randomness in Algorithm 2 comes from the step in which each EV samples from a probability distribution to update its tentative charging profile. The probability distributions are designed such that on average, the objective value decreases from the previous iteration.

Third, we provide performance guarantees for Algorithm 2. A charging profile x of all EVs is called stationary if the sequence of tentative charging profiles obtained by Algorithm 2 stops updating once the sequence reaches x. We prove that stationary charging profiles exist and provide a characterization for them. In particular, a charging profile is stationary if and only if it is a Nash equilibrium of the game in which each EV seeks to minimize the inner product of its own charging profile and the aggregate electricity demand. We also provide a uniform upper bound on the suboptimality of stationary charging profiles. The upper bound scales $O(1/n)$ as the number n of EVs increases. Hence, when the number n of EVs is large, all stationary charging profiles are nearly optimal.

We prove that the sequence of tentative charging profiles obtained by Algorithm 2 converges to stationary charging profiles. More specifically, we prove that the sequence of tentative charging profiles converges to the set of stationary charging profiles. Furthermore, in cases where each EV only has finitely many admissible charging profiles, the sequence of tentative charging profiles converges to a single but random stationary charging profile, i.e., the sequence may converge to different stationary charging profiles in different realizations. In numerical simulations, Algorithm 2 converges fast to nearly optimal charging profiles. In particular, after 10 iterations, the charging profile obtained by Algorithm 2 incurs a suboptimality less than 3%, for a wide range of EV penetration levels.

II. Problem formulation

We study the offline charging scheduling of EVs in this paper. In particular, we assume that EVs are available for negotiation at the beginning of a time horizon, on their charging profiles over the horizon. The goal is to compute an EV charging schedule that flattens the aggregate electricity demand while respecting the constraints of EVs.

The EV charging scheduling problem can be cast in a more general setting, where there are workloads to be distributed among several workers. The goal is to distribute the workloads such that each worker has more or less the same amount of workload. In the context of real applications, workloads can be software processes, Internet traffic, data to be transmitted, etc., and workers can be cpu cores, routers, communication channels, etc..

In this section, we first formulate a load balancing problem in Section II-A, and then describes its application to EV charging scheduling in Section II-B.

A. The load balancing problem

Consider the scenario where a number of loads need to be served over a time horizon. The loads can be divided into two categories: elastic loads whose service time and/or rate can be adjusted, and inelastic loads whose service time and rate are fixed and given. The load balancing problem considered in this paper seeks to schedule the service time and rate of elastic loads such that about the same aggregate service rate is required at each time.

Let T denote the length of the time horizon, and assume that the time horizon starts at 0 without loss of generality. In practice, T could be a day. The service rates of inelastic loads can be aggregated into a single process $b(t)$ where $t \in [0, T]$. Since inelastic loads can only be served at fixed time and rate, the process $b(t)$ is not adjustable. Hence, the process $b(t)$ is called base load in the literature.

Let n denote the total number of elastic loads and index them by $1, 2, \ldots, n$. For each elastic load i, let $x_i(\cdot) : [0, T] \to \mathbb{R}$ denote its service rate profile, i.e., a service of rate $x_i(t)$ is required by elastic load i at time t. The service time/rate of an elastic load is adjustable, and this can be captured by a region where the service rate profile of the elastic load can be optimized over. More specifically, for each elastic load i, there exists a region $\mathbb{X}_i$ such that the flexibility in x_i is captured by

$$x_i \in \mathbb{X}_i.$$

Every element in $\mathbb{X}_i$ is called an admissible service rate profile. Assume that each admissible service rate profile requires the same amount of total service, i.e., for each elastic load i, there exists $X_i \in \mathbb{R}$ such that for any admissible service rate profile $y_i \in \mathbb{X}_i$, one has

$$\int_0^T y_i(t)dt = X_i.$$

The quantity X_i is called the total service requirement.

The goal is to flatten the aggregate service rate profile

$$d := b + \sum_{i=1}^{n} x_i.$$

This is captured by minimizing the variance

$$V(d) := \frac{1}{T}\int_0^T \left(d(t) - \frac{1}{T}\int_0^T d(\tau)d\tau\right)^2 dt = \frac{1}{T}\int_0^T d^2(t)dt - \mu_d^2$$

of d, where the quantity

$$\mu_d := \frac{1}{T}\left(\int_0^T b(t)dt + \sum_{i=1}^{n} X_i\right)$$

is the time average of d. Note that μ_d is independent of x_i, and therefore minimizing the variance $V(d)$ of d over $x := (x_1, x_2, \ldots, x_n)$ is equivalent to minimizing the squared ℓ_2 norm

$$\|d\|_2^2 = \int_0^T d^2(t)dt$$

of d over x.

To summarize, the load balancing problem is formulated as follows.

$$\textbf{LB:} \ \min \quad \int_0^T \left(b(t) + \sum_{i=1}^{n} x_i(t)\right)^2 dt \tag{1a}$$

$$\text{s.t.} \ \ x_i \in \mathbb{X}_i, \quad i = 1, \ldots, n. \tag{1b}$$

Note that the problem LB can be nonconvex since the sets $\mathbb{X}_i$ can be nonconvex.

The following assumption is made on the constraint sets $\mathbb{X}_i$ throughout this work. To state the assumption, let $\mathcal{C}[0,T]$ denote the set of real-valued continuous functions defined on $[0,T]$, and equip $\mathcal{C}[0,T]$ with the ℓ_∞ norm $\|\cdot\|_\infty : \mathcal{C}[0,T] \mapsto \mathbb{R}_+$ defined as

$$\|f\|_\infty := \max_{t\in[0,T]} \{|f(t)|\}, \qquad \forall f \in \mathcal{C}[0,T].$$

We assume that the set $\mathbb{X}_i$ is compact throughout this paper.

Remark 1. *The algorithms and analysis developed in this paper apply to more general objective functions. In particular, the objective function in* (1a) *can be generalized to*

$$\int_0^T C_t\left(b(t) + \sum_{i=1}^n x_i(t)\right) dt \tag{1a'}$$

where the function $C_t : \mathbb{R} \mapsto \mathbb{R}$ represents the cost at time t and has a uniform upper bound β_t on its second derivatives $C_t''(\tau)$ over $\tau \in [0,T]$ for $t \in [0,T]$, i.e., there exists $\beta_t \in \mathbb{R}$ for $t \in [0,T]$ such that

$$C_t''(\tau) \le \beta_t, \qquad \forall t, \tau \in [0,T].$$

Note that if $C_t(x) = x^2$ for all $t \in [0,T]$, then (1a') *simplifies to* (1a)*. Also note that the objective of shaping the aggregate service profile d to track a target profile $G : [0,T] \mapsto \mathbb{R}$ can be captured by setting*

$$C_t(x) = (x - G(t))^2$$

for $t \in [0,T]$ in (1a')*.*

B. Application to EV charging

The EV charging scheduling problem can be considered as an application of the problem LB. In a power network, there are both electricity demand due to EVs and electricity demand due to other loads such as air conditioners, televisions, and lights. The charging of EVs can be postponed as long as EVs get fully charged by certain deadlines, and therefore EVs can be considered as elastic loads. Control of other loads including air conditioners, televisions, and lights is more evolved and outside the scope of this paper. Hence, we model those loads as inelastic.

At the beginning of a time horizon, the electricity utility negotiates with the EVs on their charging profiles over the time horizon, in order to achieve certain system level objectives, e.g., to flatten the aggregate electricity demand over the time horizon or to shape the aggregate electricity demand to track certain target profile. In this paper, we focus on the objective of flattening the aggregate electricity demand, and the results generalize straightforwardly to shaping the aggregate electricity demand to track certain target profile.

The objective of flattening the aggregate electricity demand is chosen for the following reasons. A flat electricity demand has the lowest peak and therefore induces the least burden to circuit devices. Consequently, the lifespan of circuit devices will be prolonged. Besides, the voltages in the network fluctuate as electricity demand varies. Hence, a flat electricity demand causes the least amount of fluctuation in network voltages and simplifies voltage regulation. At last, electricity generation and demand need to be balanced in almost real time since electricity

storage is limited. In practice, such balance is achieved by controlling the electricity generation of some power plants called "spinning reserves" to follow the demand. These power plants need to ramp up and down fast enough to follow the fluctuation in electricity demand, and are usually expensive. If electricity demand is flat, then fewer "spinning reserves" will be necessary and electricity supply cost can be significantly reduced.

Adopt the notations introduced in Section II-A, i.e., let $[0, T]$ denote the time horizon, let $b(t)$ denote the aggregate electricity demand of inelastic loads, index the EVs by $1, 2, \ldots, n$, and let $x_i(t)$ denote the charging profile of EV i for $i = 1, 2, \ldots, n$.

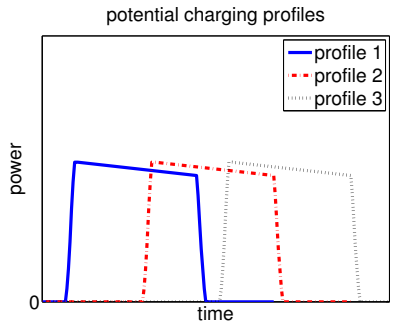

Fig. 1. Hypothetical charging profiles of an EV. They are time-shifted versions of each other.

Charging an EV at certain fixed pattern can prolong the lifespan of its battery [24]. If such patterns are enforced in the scheduling problem, then admissible charging profiles $y_i \in \mathbb{X}_i$ of an EV i are time-shifted versions of each other as illustrated in Fig. 1. In this case, the set $\mathbb{X}_i$ of admissible charging profiles is not convex and therefore the problem LB is difficult to solve, but elements y_i in $\mathbb{X}_i$ satisfy the following properties:

A1) Charging rate is uniformly bounded, i.e., there exists $\overline{x}_i \in \mathbb{R}$ such that

$$|y_i(t)| \leq \overline{x}_i, \qquad \forall t \in [0, T],\ \forall y_i \in \mathbb{X}_i.$$

A2) Ramp rate is uniformly bounded, i.e., there exists $L_i \in \mathbb{R}$ such that

$$|y_i(s) - y_i(t)| \leq L_i |s - t|, \qquad \forall s, t \in [0, T],\ \forall y_i \in \mathbb{X}_i.$$

A3) Total energy request is fixed, i.e., there exists $X_i \in \mathbb{R}$ such that

$$\int_0^T y_i(t) dt = X_i, \qquad \forall y_i \in \mathbb{X}_i.$$

A4) The ℓ_2 norm of charging profile is fixed, i.e., there exists $Y_i \in \mathbb{R}$ such that

$$\int_0^T y_i^2(t) dt = Y_i, \qquad \forall y_i \in \mathbb{X}_i.$$

Remark 2. *Let* $\mathrm{conv}(\mathbb{X}_i)$ *denote the convex hull of* $\mathbb{X}_i$*, i.e.,*

$$\mathrm{conv}(\mathbb{X}_i) = \left\{ \sum_{k=1}^m \theta_k y_k \mid m \in \mathbb{Z}^+,\ \theta_k \geq 0 \textit{ for all } k,\ \sum_{k=1}^m \theta_k = 1,\ y_k \in \mathbb{X}_i \textit{ for all } k \right\},$$

then every element $y_i \in \mathrm{conv}(\mathbb{X}_i)$ also satisfies A1, A2, and A3. Therefore, the set $\mathrm{conv}(\mathbb{X}_i)$ is uniformly bounded (by A1) and equi-continuous (by A2). The set $\mathrm{conv}(\mathbb{X}_i)$ is also closed since we assume $\mathbb{X}_i$ to be compact throughout this paper. Therefore, the set $\mathrm{conv}(\mathbb{X}_i)$ is compact in the metric space $(\mathcal{C}([0,T]), \|\cdot\|_\infty)$ by the Arzela-Ascoli Theorem [26, Theorem 11.18].

Remark 3. *The sets $\mathbb{X}_i$'s can be simplified to be finite (a compact set can be approximated by a set of finite elements to arbitrary precision) in order to ease the computation of the algorithm developed in Section III-B. More specifically, we restrict the times when an EV starts charging to finite choices. For instance, if an EV can start charging at any time from 20:00 to 3:00, then we enforce the EV to start charging only at 20:00, 20:05, ..., 3:00. Consequently, the set $\mathbb{X}_i$ of admissible charging profiles is finite.*

Let m denote the number of elements in $\mathbb{X}_i$, and denote these elements by $y_1, y_2, \ldots, y_m$. Every element $z \in \mathrm{conv}(\mathbb{X}_i)$ is a convex combination of $y_1, y_2, \ldots, y_m$, i.e., there exists $\theta = (\theta_1, \theta_2, \ldots, \theta_m)$ such that

$$\theta_k \geq 0 \text{ for all } k, \ \sum_{k=1}^{m} \theta_k = 1, \ z = \sum_{k=1}^{m} \theta_k y_k.$$

It is straightforward to show that the map

$$f : z \mapsto \theta$$

is continuous and one-to-one from $\mathrm{conv}(\mathbb{X}_i)$ to the probability simplex

$$\mathbb{P}^m = \left\{ p = (p_1, p_2, \ldots, p_m) \mid p_i \geq 0 \text{ for all } i, \ \sum_{i=1}^{m} p_i = 1 \right\},$$

i.e., $\mathrm{conv}(\mathbb{X}_i)$ is homeomorphic to $\mathbb{P}^m$.

III. Distributed algorithms

The focus of this paper is the design and analysis of a distributed algorithm to solve the problem LB. In the special case where constraint sets $\mathbb{X}_i$ are convex, we propose a distributed algorithm Algorithm 1 in Section III-A. In each iteration of Algorithm 1, a coordinator broadcasts a control signal in order to guide the loads to update their tentative service rate profiles. In the general case where constraint sets $\mathbb{X}_i$ are not necessarily convex but their convex hulls $\mathrm{conv}(\mathbb{X}_i)$ are, we propose a randomized distributed algorithm Algorithm 2 in Section III-B. Algorithm 2 shares the same coordinator operations as Algorithm 1, but differs in the load operations.

A. Special case: convex constraint sets $\mathbb{X}_i$

When $\mathbb{X}_i$ is convex for $i \in \mathcal{N}$, the problem LB can be solved by a gradient projection approach as in our prior work [20]. Here, we modify the algorithm proposed in [20] by allowing loads to update their tentative service rate profiles at different speeds in order to increase the convergence rate.

Infrastructure requirement

The algorithm assumes the availability of a coordinator that can broadcast control signals to all loads and receive feedbacks from the loads. In the application to EV charging, the coordinator can be an electricity utility or a lower level aggregator, and the communication between the EVs and the coordinator can be enabled by for example the Home Area Networks.

Algorithm statement

To state the algorithm, define the inner product $\langle \cdot, \cdot \rangle : \mathcal{C}[0,T] \times \mathcal{C}[0,T] \to \mathbb{R}$ for two continuous functions f, g on $[0,T]$ as

$$\langle f, g \rangle := \int_0^T f(t)g(t)dt,$$

and define the ℓ_2 norm $\| \cdot \| : \mathcal{C}[0,T] \to \mathbb{R}$ for a continuous function f on $[0,T]$ as

$$\|f\| := \sqrt{\langle f, f \rangle}.$$

Algorithm 1 Constraint set $\mathbb{X}_i$ is compact and convex for $i \in \mathcal{N}$

Input: Horizon length T, base load b, constraint sets $\mathbb{X}_i$ for $i \in \mathcal{N}$, componentwise strictly positive parameter $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$, tolerance $\epsilon > 0$.

Output: A service profile $x = (x_1, \ldots, x_n)$.

Initialization

Each load $i \in \mathcal{N}$ initializes $x_i^{(0)} \leftarrow 0$;

$k \leftarrow 1$;

Repeat

The coordinator collects $x_i^{(k-1)}$ from all loads and computes a control signal

$$g^{(k)} \leftarrow \frac{b + \sum_{i=1}^n x_i^{(k-1)}}{\sum_{i=1}^n c_i}; \tag{2}$$

Each load $i \in \mathcal{N}$ receives the control signal $g^{(k)}$, and updates its service rate profile as

$$x_i^{(k)} \leftarrow \underset{x_i \in \mathbb{X}_i}{\operatorname{argmin}} \; 2c_i \left\langle g^{(k)}, x_i \right\rangle + \left\| x_i - x_i^{(k-1)} \right\|^2; \tag{3}$$

$k \leftarrow k + 1$;

Until $k > 2$ and $\left\| g^{(k-1)} - g^{(k-2)} \right\| < \epsilon$.

Return $x \leftarrow x^{(k-1)}$.

The algorithm is iterative and works as follows. In each iteration $k = 1, 2, \ldots$:

- The coordinator receives tentative service profiles $x_i^{(k-1)}$ from the loads, and computes the normalized aggregate service profile $g^{(k)}$ (normalize by $\sum_{i=1}^n c_i$). Then, the coordinator broadcasts $g^{(k)}$ as a control signal to all loads in order to coordinate the loads in updating their tentative service profiles.

- Each load updates its tentative service profile by solving (3) after receiving the broadcast control signal $g^{(k)}$ from the coordinator. The objective in (3) is the sum of two terms $\left\langle g^{(k)}, x_i \right\rangle$ and $\left\| x_i - x_i^{(k-1)} \right\|^2$ weighted by $2c_i$ and 1 respectively. The first term $\left\langle g^{(k)}, x_i \right\rangle$, if interpreting $g^{(k)}$ as the time-varying service price profile, is the service cost associated with service profile x_i. The second term $\left\| x_i - x_i^{(k-1)} \right\|^2$ penalizes the deviation from x_i to the service profile $x_i^{(k-1)}$ computed in the previous iteration $k-1$. The second term is added in (3) to ensure the convergence of Algorithm 1.

The parameters c_i weigh the two terms in (3) and can be chosen appropriately to speed up the convergence rate of Algorithm 1. In particular, assume that A3 holds, i.e., the total service X_i required by a load i is a constant that does not depend on the choice of admissible service profile $y_i \in \mathbb{X}_i$ for $i \in \mathcal{N}$, and c_i is set as $c_i = X_i$ for $i \in \mathcal{N}$. The reason will become clear after Theorem 2.

Optimality

Algorithm 1 solves the problem LB in the special case where constraint sets $\mathbb{X}_i$ are convex and compact. To state the result, for two service rate profiles $x = (x_1, \ldots, x_n)$ and $x' = (x'_1, \ldots, x'_n)$, define the distance $d(x, x')$ between x and x' as

$$d(x, x') := \sqrt{\sum_{i=1}^{n} \|x_i - x'_i\|^2}.$$

Let $\mathbb{O}^*$ denote the set of optimal solutions of LB, and define the distance $d(x, \mathbb{O}^*)$ from a service rate profile x to the set $\mathbb{O}^*$ as

$$\mathrm{dist}(x, \mathbb{O}^*) := \inf_{x^* \in \mathbb{O}^*} d(x, x^*).$$

If $\mathrm{dist}(x, \mathbb{O}^*) = 0$, then x is in the closure $\overline{\mathbb{O}^*}$ of $\mathbb{O}^*$. If $\mathrm{dist}(x, \mathbb{O}^*)$ is small, then x is "close" to $\mathbb{O}^*$.

Theorem 1. *If $\mathbb{X}_i$ is compact and convex for $i \in \mathcal{N}$, then the sequence $\{x^{(0)}, x^{(1)}, \ldots, x^{(k)}, \ldots\}$ of service profiles computed in Algorithm 1 converges to the set $\mathbb{O}^*$ of optimal solutions of LB as $k \to \infty$, i.e.,*

$$\mathrm{dist}(x^{(k)}, \mathbb{O}^*) \to 0 \text{ as } k \to \infty.$$

Theorem 1 implies that Algorithm 1 obtains an optimal service profile. It is proved in Appendix A. Note that no further assumption is made on the parameter c in Theorem 1, i.e., Algorithm 1 eventually solves the problem LB for any parameter c that is componentwise strictly positive. However, a good choice of c can significantly improve the convergence rate of Algorithm 1.

Convergence rate

The parameter c is picked for a faster convergence of Algorithm 1, and its choice is based on the following observations. First, c_i affects how fast a load i updates its tentative service profile x_i: with a larger c_i, more weight is placed on the first summand in (3) and therefore the second summand in (3), a penalty term on service profile updates, is less impactful. Consequently, load i tends to make larger updates in x_i.

Second, the second summand in (3) scales quadratically with total service X_i, while $\left\langle g^{(k)}, x_i \right\rangle$ in the first summand scales linearly with total service X_i. Hence, if c_i is set proportional to X_i, then the two summands in (3)

are scale invariant, i.e., the two summands in (3) remain the same ratio as X_i varies over a wide range. According to the first observation, this choice of c_i enforces loads with larger total service X_i to take larger steps in their service profile updating, such that the tentative service profiles of different loads converge to the final service profiles at a "uniform" rate.

The following pathological example explains the benefit of setting c_i proportional to X_i.

Example 1. *Consider two instances I1 and I2 of the problem LB. In Instance I1, there are 3 identical loads 0, 1, 2, each with constraint set $\mathbb{X}$ and total service X. It is not difficult to verify that if we set $c_0 = c_1 = c_2$, then Algorithm 1 obtains an optimal service profile within 1 iteration.*

In Instance I2, there are only two loads 1' and 2'. While load 2' is the same as load 2, i.e., load 2' has constraint set $\mathbb{X}$ and requests total service X, load 1' is effectively the aggregate of load 0 and load 1, i.e., load 1' has constraint set $\mathbb{X}_{1'} = 2\mathbb{X}$ and requests total service $X_{1'} = 2X$. If we still set $c_{1'} = c_{2'}$, then Algorithm 1 no longer converges within 1 iteration. This is because load 1' has a larger feasible set $\mathbb{X}_{1'}$ to explore than load 2', but can only update its tentative service profiles $x_{1'}$ at the same speed as load 2'. To remedy this, one can set $c_{1'} = 2c_{2'}$. Interestingly, with this choice of $c_{i'}$, one can prove that $x_{1'}^{(k)} = 2x_0^{(k)} = 2x_1^{(k)}$ and $x_{2'}^{(k)} = x_2^{(k)}$ for $k = 0, 1, 2, \ldots$, i.e., Algorithm 1 behaves equivalently in Instances I1 and I2 in the sense that load $1'$ is effectively the aggregate of load 0 and load 1. Hence, Algorithm 1 still converges within 1 iteration, highlighting the benefit of setting c_i proportional to X_i.

We can generalize Example 1 to multiple-load setups as stated in the following theorem.

Theorem 2. *Assume that $\mathbb{X}_i$ is compact and convex, and that $c_i = X_i > 0$ for each i. Consider the following two instances of the problem LB.*

I1) *There are $n+1$ loads $0, 1, \ldots, n$, and load 0 and load 1 are identical.*

I2) *There are n loads $1', 2', \ldots, n'$ where load 1' is the aggregate of load 0 and load 1 while loads $2' \sim n'$ are the same as loads $2 \sim n$, i.e.,*

$$\mathbb{X}_{1'} = 2\mathbb{X}_0 = 2\mathbb{X}_1, \quad \mathbb{X}_{i'} = \mathbb{X}_i \text{ for } i = 2, 3, \ldots, n.$$

The sequence $\left\{[x_i^{(0)}]_{i=0}^n, [x_i^{(1)}]_{i=0}^n, \ldots, [x_i^{(k)}]_{i=0}^n, \ldots\right\}$ of tentative service profiles computed in Algorithm 1 in Instance I1 is equivalent to the sequence $\left\{[x_{i'}^{(0)}]_{i=1}^n, [x_{i'}^{(1)}]_{i=1}^n, \ldots, [x_{i'}^{(k)}]_{i=1}^n, \ldots\right\}$ of tentative service profiles computed in Algorithm 1 in Instance I2 in the sense that

$$x_{1'}^{(k)} = 2x_0^{(k)} = 2x_1^{(k)}, \qquad x_{i'}^{(k)} = x_i^{(k)} \text{ for } i = 2, 3, \ldots, n \tag{4}$$

for $k = 0, 1, 2, \ldots$.

Theorem 2 implies that the aggregation of identical loads into a bigger load does not change the convergence rate of Algorithm 1 provided that $c_i = X_i$ for all i. The theorem is proved in Appendix B. Hence, we set $c_i = X_i$ for all i in Algorithm 1.

Connection to prior work

Algorithm 1 generalizes the distributed algorithm proposed in [20] by introducing a parameter c to regulate the tentative service profile updates of different loads. In particular, if $c_i = 1$ for all i, then Algorithm 1 reduces to the algorithm proposed in [20]. The freedom of choosing c in Algorithm 1 can be exploited to improve the convergence rate as exemplified in Example 1. In this paper, we set c as $c_i = X_i$ for all i.

B. General case: nonconvex constraint sets $\mathbb{X}_i$

The problem LB can be solved by Algorithm 1 in the special case where constraint sets $\mathbb{X}_i$ are convex. However, the constraint sets $\mathbb{X}_i$ are not convex for many applications including the application to EV charging, and therefore Algorithm 1 does not directly apply. In this section, we seek to modify Algorithm 1 such that certain types of nonconvex constraint sets $\mathbb{X}_i$ can be dealt with.

Motivated by the application to EV charging, we assume that each $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, and A4 in the rest of this section. It follows from Remark 2 that $\mathrm{conv}(\mathbb{X}_i)$ is compact for each i.

Infrastructure requirement

We aim at an algorithm with the same infrastructure requirement and information flow pattern as Algorithm 1 so that loads with convex and nonconvex constraint sets can be scheduled simultaneously. In particular, we assume the availability of a coordinator that is capable of broadcasting control signals to all loads and receiving feedbacks.

The information flow pattern of Algorithm 1 (as well as the algorithm to be designed) is described in Fig. 2 where a coordinator and multiple loads exchange information in multiple iterations to agree on a service schedule $x = (x_1, x_2, \ldots, x_n)$. In each iteration k, the coordinator computes F to obtain a control signal $g^{(k)}$ based on the base load profile b and tentative service profiles $x_1^{(k-1)}, \ldots, x_n^{(k-1)}$ computed by loads in the previous iteration, and each load i computes G to update its tentative service profile $x_i^{(k)}$ based on the control signal $g^{(k)}$, the tentative service profile $x_i^{(k-1)}$ it computes in the previous iteration, and its set of admissible service profiles $\mathbb{X}_i$. Here F and G are the functions to be designed.

In order to schedule the loads with convex and nonconvex constraints simultaneously, the algorithm to be designed must share the same coordinator operations, i.e., the same function F, as Algorithm 1. However, the load operations, i.e., the function G, can be different for the algorithm to be designed and Algorithm 1.

When there are both loads with convex constraints and loads with nonconvex constraints in the system, the coordinator does not need to behave differently and simply computes F in each iteration, while each load, depending on whether it has convex constraints or nonconvex constraints, computes the G for Algorithm 1 or the G for the algorithm to be designed, respectively.

Failure of deterministic algorithms

Any algorithm with the information flow pattern in Fig. 2 and a deterministic function G either diverges or converges to synchronized solutions in the setting where all loads are identical.

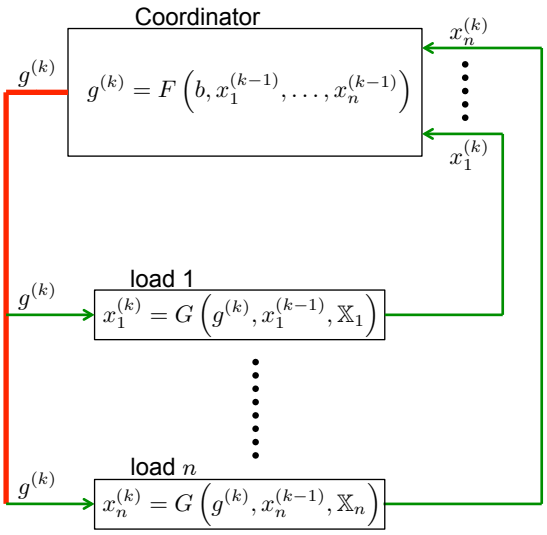

Fig. 2. Information flow pattern in the proposed iterative, distributed decision-making process. The coordinator knows the base load b and each load i knows its admissible service profiles $\mathbb{X}_i$.

Proposition 1 (Proposition 1, [27])**.** *If $\mathbb{X}_1 = \mathbb{X}_2 = \ldots = \mathbb{X}_n$ and G is a deterministic function, then any algorithm with the information flow pattern in Fig. 2 obtains synchronized service profiles in all iterations, i.e.,*

$$x_1^{(k)} = x_2^{(k)} = \cdots = x_n^{(k)}, \qquad k \geq 0.$$

Proposition 1 implies that if loads are identical, then deterministic algorithms obtain synchronized service profiles. When $\mathbb{X}_i$'s are convex, synchronized service profiles can still be optimal and that's why Algorithm 1, a deterministic algorithm, can be optimal in that setting. However when $\mathbb{X}_i$'s are nonconvex, even the best synchronized service profile can be (and are usually) far from optimal. This highlights the necessity of randomized algorithms.

Remark 4. *Proposition 1 only discusses deterministic algorithms using the information flow pattern in Fig. 2. It is possible to find a deterministic algorithm that flattens the aggregate load using other information flow patterns. For instance, consider the following information flow pattern. In each iteration k, the coordinator uses the most recently calculated $x_1, \ldots, x_n$ to compute the broadcast signal g, and only one load (in turn) updates its tentative service profile. With this information flow pattern, there exists deterministic algorithms that flatten the aggregate load, however, at a much higher communication overhead and a much longer delay.*

Idea of randomized algorithms

A randomized algorithm adopts a randomized function G and works as follows. In each iteration k, each load i first computes a probability distribution $P_i^{(k)}$ over its admissible service profiles $\mathbb{X}_i$ and then samples from the probability distribution $P_i^{(k)}$ to update its tentative service profile $x_i^{(k)}$. The algorithmic challenge is the design and computation of a probability distribution $P_i^{(k)}$ that ensures the convergence of a randomized algorithm.

The probability distributions $P_i^{(k)}$ are designed such that the objective values in consecutive iterations of the algorithm decrease in expectation. With this property and lower boundedness of the objective values, there is a prominent theorem saying that the objective values almost surely converge as iterations continue.

Algorithm statement

To state the algorithm, let $\Theta(A)$ denote the set of probability distributions over a set A. For example, if A has m elements, then any probability distribution P over A can be described by a point $p = (p_1, p_2, \ldots, p_m)$ in the probability simplex $\mathbb{P}^m$, i.e., there exists a one-to-one mapping between $\Theta(A)$ and $\mathbb{P}^m$. Given $f : A \to \mathbb{R}$ and a probability distribution P over A, let

$$\mathbb{E}_P[f] := \int_A f(y) dP$$

denote the expectation of f with respect to the probability distribution P.

Algorithm 2 Constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$

Input: Horizon length T, base load b, constraint sets $\mathbb{X}_i$ for $i \in \mathcal{N}$, componentwise strictly positive parameter $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$, tolerance $\epsilon > 0$.

Output: A service profile $x = (x_1, \ldots, x_n)$.

Initialization

Each load $i \in \mathcal{N}$ initializes $x_i^{(0)} \leftarrow 0$;

$k \leftarrow 1$;

Repeat

The coordinator collects $x_i^{(k-1)}$ from all loads and computes a control signal

$$g^{(k)} \leftarrow \frac{b + \sum_{i=1}^n x_i^{(k-1)}}{\sum_{i=1}^n c_i}; \tag{5}$$

Each load $i \in \mathcal{N}$ receives the control signal $g^{(k)}$, computes a probability distribution $P_i^{(k)}$ over $\mathbb{X}_i$ as

$$P_i^{(k)} \leftarrow \operatorname*{argmin}_{P_i \in \Theta(\mathbb{X}_i)} 2c_i \left\langle \frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j \neq i} c_j},\ \mathbb{E}_{P_i}[x_i] \right\rangle + \left\| \mathbb{E}_{P_i}[x_i] - x_i^{(k-1)} \right\|^2, \tag{6}$$

and samples from $P_i^{(k)}$ to update a tentative service profile $x_i^{(k)} \in \mathbb{X}_i$;

$k \leftarrow k + 1$;

Until $k > 2$ and $\left\| g^{(k-1)} - g^{(k-2)} \right\| < \epsilon$.

Return $x \leftarrow x^{(k-1)}$.

Algorithm 2 is similar to Algorithm 1, but with different load operations. In each iteration k of Algorithm 2, a load i does not directly compute an $x_i^{(k)} \in \mathbb{X}_i$ as in Algorithm 1, but first computes a probability distribution $P_i^{(k)}$ over $\mathbb{X}_i$ by solving (6) and then samples from $P_i^{(k)}$ to update its tentative service profile $x_i^{(k)}$. The objective in (6) is chosen such that the objective values evaluated in consecutive iterations of Algorithm 2 decrease in expectation.

The objective in (6) differs from that in (3) in two aspects. First, the rate profile $x_i \in \mathbb{X}_i$ in (3) is replaced by the expectation $\mathbb{E}_{P_i}[x_i]$ of a probability distribution P_i over $\mathbb{X}_i$ in (6). Second, the first summand in the objective of

(6) has a negative term $-x_i^{(k-1)}$ and a multiplicative factor $\eta := \sum_{j=1}^n c_j / \sum_{j\neq i} c_j$ in front of $g^{(k)}$ and . Note that $x_i^{(k-1)}$ is negligible in comparison with $b + \sum_{j=1}^n x_j^{(k-1)} = g^{(k)} \sum_{j=1}^n c_j$, so the negative term $-x_i^{(k-1)} / \sum_{j\neq i} c_j$ is negligible. Besides, if there is no single c_i that dominates the sum $\sum_{j=1}^n c_j$, then the multiplicative factor $\sum_{j=1}^n c_j / \sum_{j\neq i} c_j$ in front of $g^{(k)}$ is approximately 1. To summarize,

$$\frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j\neq i} c_j} \approx g^{(k)}, \tag{7}$$

i.e., the objective in (6) is approximately equal to the objective in (3) with x_i substituted by $\mathbb{E}_{P_i}[x_i]$.

The computation of a probability distribution $P_i^{(k)}$ can be difficult unless a simple representation of $P_i^{(k)}$ is available. As discussed in Remark 3, in the application to EV charging, the set $\mathbb{X}_i$ can be reduced to a finite set in order to ease the computation of $P_i^{(k)}$. In this case, $P_i^{(k)}$ can be represented by its probability mass function over $\mathbb{X}_i$, i.e., let m denote the number of elements in $\mathbb{X}_i$, then there is a one-to-one mapping between $\Theta(\mathbb{X}_i)$ (where $P_i^{(k)}$ takes values in) and the probability simplex $\mathbb{P}^m$. Moreover, the map $f : \Theta(\mathbb{X}_i) \mapsto \mathrm{conv}(\mathbb{X}_i)$ defined by

$$f(P_i) = \mathbb{E}_{P_i}[x_i]$$

is one-to-one. Hence, computing a $P_i^{(k)} \in \Theta(\mathbb{X}_i)$ is equivalent to finding an $\mathbb{E}_{P_i^{(k)}}[x_i] \in \mathrm{conv}(\mathbb{X}_i)$.

Remark 5. *Load i needs to know $C := \sum_{j=1}^n c_j$ in order to compute* (6)*, and the coordinator can broadcast C together with $g^{(k)}$ to provide this piece of information. Alternatively, load i can also use the approximation in* (7) *so that it no longer needs to know C. In this case, the objective load i minimizes is the same as the objective in* (3) *with x_i substituted by $\mathbb{E}_{P_i}[x_i]$, i.e., each load, whether it has convex or nonconvex constraint, minimizes the same objective.*

Stationary service profiles and optimality

We will show that the sequence of tentative service profiles obtained by Algorithm 2 converges to certain nearly optimal service profiles. To start with, we characterize these service profiles and show that they are nearly optimal.

Definition 1. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. A service profile $x = (x_1, \ldots, x_n)$ is stationary for Algorithm 2, if $x^{(k-1)} = x$ implies $x^{(k)} = x$ with probability 1 for $k \geq 1$.*

That is, stationary service profiles are those once hit, tentative service profiles stop updating. Note that stationary service profiles do not exist for many randomized algorithms including simulated annealing [28] and genetic algorithm [29]. In these two widely known algorithms, the probability

$$\mathbb{P}\{x^{(k)} \neq x \mid x^{(k-1)} = x\} > 0 \tag{8}$$

is nonzero for any $k \geq 1$ and any x, and therefore stationary service profiles do not exist for these two algorithms.

The reason simulated annealing and genetic algorithm satisfy (8) is that they seek to escape from local optimal points, and this objective is achieved at the cost of escaping from global optimal points as well. The methodology is that, the probability of obtaining a global optimal point increases via a so-called "cooling"/"evolving" process

as iterations continue. Though the cooling/evolving process is usually very slow, eventually a global optimal point is obtained with a strictly positive probability.

Algorithm 2 seeks to find a nearly optimal service profile within few iterations, instead of having a slow cooling (evolving) process to find an optimal service profile (which is NP hard). At the cost of small suboptimality (Theorem 4), Algorithm 2 has stationary service profiles (Theorem 3) and converges to these stationary service profiles quickly—we stop Algorithm 2 after the first 20 iterations in the simulations provided in Section IV.

A characterization of the stationary service profiles of Algorithm 2 is provided in the following theorem.

Theorem 3. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. A service profile $x = (x_1, \ldots, x_n)$ is stationary for Algorithm 2 if and only if it is a Nash-equilibrium of the game in which each load $i \in \mathcal{N}$ seeks to minimize $\left\langle b + \sum_{j=1}^n x_j, x_i \right\rangle$ over $x_i \in \mathbb{X}_i$.*

Theorem 3 implies that the stationary service profiles of Algorithm 2 happen to be the Nash-equilibria of a game. It is proved in Appendix C. In this game, each load $i \in \mathcal{N}$ aims to minimize the inner product of its own service profile x_i and the aggregate service profile of all loads $b + \sum_{j=1}^n x_j$. Hence at a Nash equilibrium, load i chooses to be serviced at times when the aggregate load $b + \sum_{j=1}^n x_j$ is low. Therefore, stationary service profiles are likely to be close to flat. Indeed, the following theorem provides an upper bound on the suboptimality of stationary service profiles of Algorithm 2. Before stating the theorem, recall the definition of Y_i in A4.

Theorem 4. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. Let x^s be a stationary service profile for Algorithm 2 and x^* be a global optimum of the problem LB, then*

$$\left\| b + \sum_{i=1}^n x_i^s \right\|^2 - \left\| b + \sum_{i=1}^n x_i^* \right\|^2 \leq 4 \sum_{i=1}^n Y_i. \tag{9}$$

If further, elements in $\mathbb{X}_i$ are nonnegative, i.e., $y_i \geq 0$ componentwise for all $i \in \mathcal{N}$ and all $y_i \in \mathbb{X}_i$, then

$$\left\| b + \sum_{i=1}^n x_i^s \right\|^2 - \left\| b + \sum_{i=1}^n x_i^* \right\|^2 \leq 2 \sum_{i=1}^n Y_i. \tag{10}$$

Theorem 4 implies that the suboptimality of a stationary service profile is uniformly upper bounded by constant $4 \sum_{i=1}^n Y_i$, which scales $O(n)$ as the number n of loads increases. The theorem is proved in Appendix D. Noting that the optimal objective value scales $O(n^2)$ as n increases, the suboptimality ratio defined as

$$\mathrm{SubOpt}(x^s) := \frac{\|b + \sum_{i=1}^n x_i^s\|^2 - \|b + \sum_{i=1}^n x_i^*\|^2}{\|b + \sum_{i=1}^n x_i^*\|^2} \leq \frac{4 \sum_{i=1}^n Y_i}{\|b + \sum_{i=1}^n x_i^*\|^2} \tag{11}$$

has an upper bound that scales $O(1/n)$ as n increases. Hence, all stationary service profiles of Algorithm 2 are close to optimal when n is sufficiently large. When $\mathbb{X}_i$'s are composed of nonnegative elements, the suboptimality ratio upper bound can be improved to

$$\mathrm{SubOpt}(x^s) \leq \frac{2 \sum_{i=1}^n Y_i}{\|b + \sum_{i=1}^n x_i^*\|^2}. \tag{12}$$

Martingale and convergence

Algorithm 2 is designed such that the objective values obtained in consecutive iterations decrease in expectation, since then there is a martingale theory ensuring the convergence of objective values. To state the result, let

$$L_k := \left\| b + \sum_{i=1}^{n} x_i^{(k)} \right\|^2$$

denote the objective value obtained in iteration k of Algorithm 2 for $k \geq 1$.

Proposition 2. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. The stochastic process $\{L_k : k \geq 1\}$ is a supermartingale with respect to the process $\{x^{(k)} : k \geq 1\}$, i.e.,*

$$\mathbb{E}\left[L_k \mid x^{(k-1)}, x^{(k-2)}, \ldots, x^{(0)} \right] \leq L_{k-1}, \quad k = 2, 3, \ldots.$$

Proposition 2 implies that $\{L_k : k \geq 1\}$ is a nonnegative supermartingale. It is proved in Appendix E. Martingale theory tells us that a nonnegative supermartingale almost surely converges [25], as stated in the following corollary.

Corollary 1. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. The sequence $L_0, L_1, \ldots, L_k, \ldots$ converges almost surely as $k \to \infty$.*

In fact, the objective in (6) is designed such that $\{L_k : k \geq 1\}$ forms a nonnegative supermartingale and therefore convergences. However, the convergence of objective values does not imply the convergence of tentative service profiles, though the reverse is true. Hence, we prove the following stronger result regarding the convergence of tentative service profiles. To state the result, let $\mathbb{S}$ denote the set of stationary service profiles of Algorithm 2.

Theorem 5. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. The set $\mathbb{S}$ of stationary service profiles of Algorithm 2 is nonempty, and the sequence $x^{(0)}, x^{(1)}, \ldots, x^{(k)}, \ldots$ of tentative service profiles computed in Algorithm 2 converges almost surely to the set $\mathbb{S}$ as k tends to infinity, i.e.,*

$$\operatorname{dist}(x^{(k)}, \mathbb{S}) \xrightarrow{\text{a.s.}} 0 \text{ as } k \to \infty.$$

Theorem 5 implies that the sequence of tentative service profiles computed in Algorithm 2 converges to the set of stationary service profiles. The theorem is proved in Appendix F. However, convergence to a set does not imply convergence to a point since there can be for example limit cycles [30, Chapter 7.0]. In the special case where $\mathbb{X}_i$'s have finitely many elements, convergence to a point can be established as in the following Corollary.

Corollary 2. *Assume that constraint set $\mathbb{X}_i$ is compact and satisfies A1, A2, A3, A4 for $i \in \mathcal{N}$. If $\mathbb{X}_i$ has finitely many elements for $i \in \mathcal{N}$, then the sequence $x^{(0)}, x^{(1)}, \ldots, x^{(k)}, \ldots$ of tentative service profiles computed in Algorithm 2 converges almost surely to a random service profile x^∞ that takes values in $\mathbb{S}$, i.e.,*

$$x^{(k)} \xrightarrow{\text{a.s.}} x^\infty \text{ as } k \to \infty$$

where x^∞ is a random variable that takes values in $\mathbb{S}$.

Corollary 2 implies that the sequence of tentative service profiles computed in Algorithm 2 converges to a random stationary service profile. The corollary is proved in Appendix G. In the application to EV charging, the sets $\mathbb{X}_i$'s are finite, and therefore Algorithm 2 converges almost surely to a single (random) stationary service profile.

IV. Case Studies

We evaluate Algorithm 2 numerically in the application to EV charging in this section. In particular, we investigate how fast Algorithm 2 converges, and how close to optimal are the charging profiles obtained by Algorithm 2 after a certain number of iterations.

The scheduling horizon T is considered to be 24 hours, and discretized into 96 slots, each of 15 minutes. The base load is set to match the average residential load profile in the service area of Southern California Edison from 20:00 on 02/13/2011 to 20:00 on 02/14/2011 [31]. We consider different penetration levels of EVs. For simplicity, we assume that each EV is charged according to the pattern in Fig. 3—it consumes 3.3kW power consecutively for 4 hours [23]. We also assume that each EV can start charging at 20:00, 20:15, ..., 16:00—as long as it can finish the 4-hour charging before the end of the time horizon.

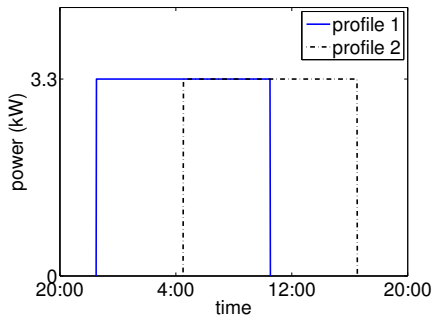

Fig. 3. Two examples of admissible charging profiles of an EV. The EV needs to charge its battery at 3.3kW power consecutively for 4 hours.

A. Convergence rate

Recall that $x^{(k)}$ denotes the tentative charging profile computed in iteration k of Algorithm 2, and note that the stochastic process $\{x^{(k)} : k \geq 1\}$ is a Markov chain [25, Chapter 6]. Define *escape probabilities*

$$P_{\text{escape}}^{(k)} := \mathbb{P}\{x^{(k)} \neq x^{(k-1)} \mid x^{(k-1)}\}, \quad k \geq 1.$$

If $P_{\text{escape}}^{(k)} = 0$, then $x^{(k-1)}$ is a stationary charging profile; otherwise, it is straightforward to verify that $1/P_{\text{escape}}^{(k)}$ is the expected number of iterations it takes before tentative charging profile gets updated. For example, if $P_{\text{escape}}^{(k)} = 0.5$, then on average a tentative charging profile update happens after 2 iterations. At such an updating speed, we may want to stop the iterations in Algorithm 2.

To visualize the relationship between tentative charging profile updates and escape probabilities, we show the average aggregate load profiles (per household) and the escape probabilities in the first 20 iterations of Algorithm

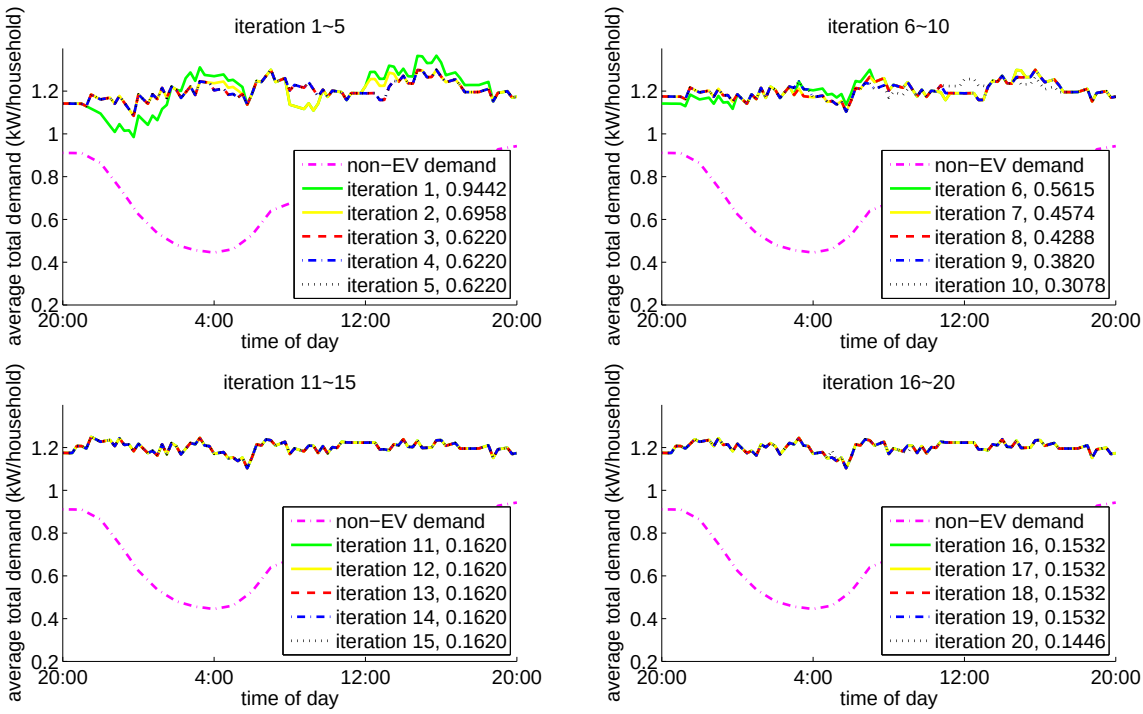

Fig. 4. Average aggregate load profiles (per household) and escape probabilities in the first 20 iterations of Algorithm 2 in a 100% EV penetration case. Escape probabilities for different iterations are shown in the legends.

2 in a 100% EV penetration case (one EV per household) in Figure 4. It can be seen that the aggregate load profile only gets slightly updated from iteration 6 to iteration 10 where the escape probability ranges from 0.3 to 0.6; and that the aggregate load profile is updated only twice from iteration 11 to iteration 20 where the escape probability is below 0.2. To summarize, the escape probability $P_{\text{escape}}^{(k)}$ measures "closeness" to stationary charging profiles: the smaller $P_{\text{escape}}^{(k)}$, the closer $x^{(k-1)}$ is to a stationary charging profile. In particular, if $P_{\text{escape}}^{(k)} < 0.5$, then $x^{(k-1)}$ is nearly stationary.

Fig. 5 shows the (average) escape probability $P_{\text{escape}}^{(k)}$ (of 10 simulations) in the first 40 iterations of Algorithm 2 at different EV penetration levels. It can be seen that the average escape probability goes below 0.5 within 20 iterations for all penetration levels. Hence, we stop Algorithm 2 after 20 iterations. We call $x^{(20)}$ the output charging profile of Algorithm 2 hereafter.

B. Suboptimality ratio

Fig. 6 shows the average aggregate load profile in iteration 20 of Algorithm 2 at different EV penetration levels. It can be seen that the average aggregate load profile is close to flat even with only 20% EV penetration. Note that completely flat aggregate load profile is not achievable since the charging rate of an EV is either 0 or 3.3kW.

Now we theoretically quantify the suboptimality ratio $\text{SubOpt}\left(x^{(20)}\right)$ [defined in (11)] of output charging profile

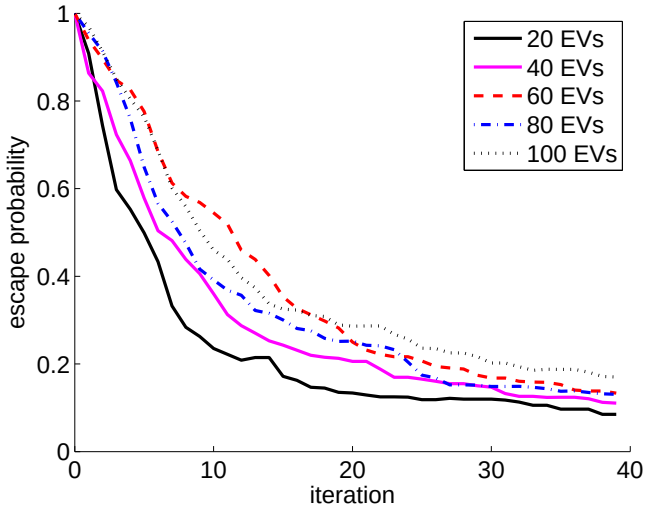

Fig. 5. Average escape probability $P_{\text{escape}}^{(k)}$ at different EV penetration levels.

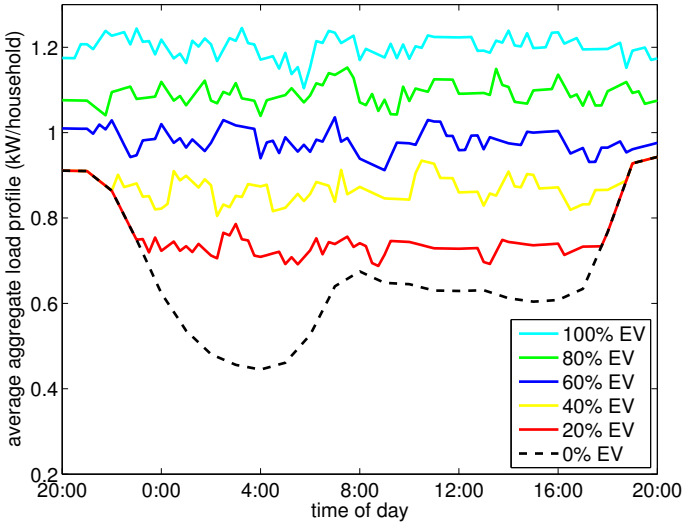

Fig. 6. Average total demand per household in iteration 20 of Algorithm DSC, with various number of EVs.

$x^{(20)}$. Since tentative charging profile updates become negligible after 20 iterations, we think of $x^{(20)}$ as a stationary charging profile and apply the suboptimality ratio upper bound derived in (12), i.e.,

$$\text{SubOpt}\left(x^{(20)}\right) \leq \frac{2\sum_{i=1}^{n} Y_i}{\|b+\sum_{i=1}^{n} x_i^*\|^2} := \text{SubOptBound} \tag{13}$$

where x^* denotes the optimal charging profile.

When the number n of EVs is small, base load b dominates aggregate EV load $\sum_{i=1}^{n} x_i^*$ and therefore

$$2\sum_{i=1}^{n} Y_i \;=\; 2\sum_{i=1}^{n} \|x_i^*\|^2 \;\leq\; 2\left\|\sum_{i=1}^{n} x_i^*\right\|^2 \;\ll\; \|b\|^2 \;\leq\; \left\|b+\sum_{i=1}^{n} x_i^*\right\|^2 .$$

Hence, $\mathrm{SubOptBound}$ is much smaller than 1 and the suboptimality ratio $\mathrm{SubOpt}\left(x^{(20)}\right)$ is small. When the number n of EVs is big, base load b is dominated by aggregate EV load and therefore

$$\mathrm{SubOptBound} \;\leq\; \frac{2\sum_{i=1}^{n} Y_i}{\left\|\sum_{i=1}^{n} x_i^*\right\|^2}. \tag{14}$$

Noting that $2\sum_{i=1}^{n} Y_i$ scales $O(n)$ as n increases and that $\left\|\sum_{i=1}^{n} x_i^*\right\|^2$ scales $O(n^2)$ as n increases, the bound $\mathrm{SubOptBound}$ scales $O(1/n)$ as n increases. Hence, the suboptimality ratio $\mathrm{SubOpt}\left(x^{(20)}\right)$ remains small when n is big.

The suboptimality upper bound $\mathrm{SubOptBound}$ for different penetration levels of EV is shown in Fig. 7. It can be seen that $\mathrm{SubOptBound}_1 \leq 2.6\%$ at all EV penetration levels.

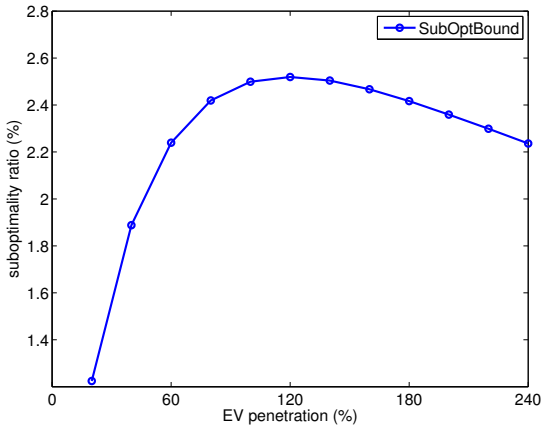

Fig. 7. Suboptimality upper bound $\mathrm{SubOptBound}$ for different EV penetration levels.

V. Conclusions

We have proposed a distributed EV charging scheduling algorithm to shape the aggregate load, that allows for a class of nonconvex constraints on admissible EV charging profiles. In particular, the algorithm applies to the case where each EV only has finitely many admissible charging profiles. The algorithm is iterative, and assumes the availability of a coordinator that can broadcast control signals to all EVs and receive feedbacks. In each iteration, the coordinator collects tentative charging profiles computed by the EVs in the previous iteration and broadcasts the normalized aggregate load profile. After receiving the broadcast signal, each EV updates its tentative charging profile by first computing a probability distribution over its admissible charging profiles, and then samples from the probability distribution to update its tentative charging profile.

The probability distribution an EV computes is designed such that the objective values in consecutive iterations of the algorithm decreases in expectation, and therefore the martingale theory ensures that the objective value converges almost surely as iterations continue. Besides, the algorithm shares the same information flow pattern as an algorithm that is designed to schedule EVs with convex constraints, and therefore both algorithms can run simultaneously in scenarios where there are both EVs with convex constraints and EVs with nonconvex constraints.

We have proved that the algorithm converges almost surely to a random stationary charging profile, and that the suboptimality ratio of the stationary charging profile has an upper bound that scales $O(1/n)$ as the number n of EVs increases. Case studies confirm that the algorithm converges fast (can stop after 20 iterations), with suboptimality ratio below 2.6% at all EV penetration levels.

REFERENCES

[1] Department of Energy, "All-electric vehicles," online at https://www.fueleconomy.gov/feg/evtech.shtml, 2014.

[2] ——, "Federal tax credits for electric vehicles purchased in or after 2010," online at https://www.fueleconomy.gov/feg/taxevb.shtml, 2014.

[3] ——, "All-electric vehicles: Compare side-by-side," online at http://www.fueleconomy.gov/feg/evsbs.shtml, 2014.

[4] Wikipedia, "Plug-in electric vehicles in the united states," online at http://en.wikipedia.org/wiki/Plug-in_electric_vehicles_in_the_United_States, 2014.

[5] GAS2, "Bmw wants to build 100000 evs annually by 2020," online at http://gas2.org/2014/03/21/bmw-wants-build-100000-evs-annually-2020/, 2014.

[6] The Wall Street Journal, "Nissan may hit electric-car sales target before 2020," online at http://online.wsj.com/news/articles/SB10001424052702303546204579440652848502402, 2014.

[7] Autobloggreen, "Tesla investor says selling 500,000 evs in 2020 is totally doable," online at http://green.autoblog.com/2014/03/11/tesla-investor-says-selling-500000-evs-in-2020-possible/, 2014.

[8] H. Streng, "Mindboggling consequences in wake of battery price drops," online at http://cleantechnica.com/2013/12/27/mondboggling-consequences-wake-battery-price-drops/, 2013.

[9] L. Kelly, A. Rowe, and P. Wild, "Analyzing the impacts of plug-in electric vehicles on distribution networks in british columbia," in *Electrical Power & Energy Conference*, 2009, pp. 1–6.

[10] C. Roe, J. Meisel, A. Meliopoulos, F. Evangelos, and T. Overbye, "Power system level impacts of phevs," in *Hawaii International Conference on System Sciences*, 2009, pp. 1–10.

[11] K. Clement, E. Haesen, and J. Driesen, "Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids," in *IEEE/PES Power Systems Conference and Exposition*, 2009, pp. 1–7.

[12] C. Quinn, D. Zimmerle, and T. H. Bradley, "The effect of communication architecture on the availability, reliability, and economics of plug-in hybrid electric vehicle-to-grid ancillary services," *Journal of Power Sources*, vol. 195, no. 5, pp. 1500–1509, 2010.

[13] J. P. Lopes, F. J. Soares, P. Almeida, and M. M. da Silva, "Smart charging strategies for electric vehicles: Enhancing grid performance and maximizing the use of variable renewable energy resources," in *Intenational Battery, Hybrid and Fuell Cell Electric Vehicle Symposium*, 2009.

[14] L. Gan, A. Wierman, U. Topcu, N. Chen, and S. H. Low, "Real-time deferrable load control: handling the uncertainties of renewable generation," in *International Conference on Future Energy Systems*, 2013, pp. 113–124.

[15] R. A. Verzijlbergh, M. O. Grond, Z. Lukszo, J. G. Slootweg, and M. D. Ilic, "Network impacts and cost savings of controlled ev charging," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1203–1212, 2012.

[16] K. Qian, C. Zhou, M. Allan, and Y. Yuan, "Modeling of load demand due to ev battery charging in distribution systems," *IEEE Transactions on Power Systems*, vol. 26, no. 2, pp. 802–810, 2011.

[17] O. Sundstrom and C. Binding, "Planning electric-drive vehicle charging under constrained grid conditions," in *International Conference on Power System Technology*, 2010, pp. 1–6.

[18] P. Richardson, D. Flynn, and A. Keane, "Optimal charging of electric vehicles in low-voltage distribution systems," *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 268–279, 2012.

[19] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2013.

[20] L. Gan, U. Topcu, and S. Low, "Optimal decentralized protocol for electric vehicle charging," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 940–951, 2013.

[21] M. Caramanis and J. M. Foster, "Management of electric vehicle charging to mitigate renewable generation intermittency and distribution network congestion," in *IEEE Conference on Decision and Control*, 2009, pp. 4717–4722.

[22] N. Rotering and M. Ilic, "Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets," *IEEE Transactions on Power Systems*, vol. 26, no. 3, pp. 1021–1029, 2011.

[23] A. Ipakchi and F. Albuyeh, "Grid of the future," *IEEE Power and Energy Magazine*, vol. 7, no. 2, pp. 52–62, 2009.

[24] L. James and L. John, "Electric vehicle technology explained," 2003.

[25] G. Grimmett and D. Stirzaker, *Probability and Random Processes*. Oxford university press, 2001.

[26] N. L. Carothers, *Real Analysis*. Cambridge University Press, 2000.

[27] L. Gan, U. Topcu, and S. H. Low, "Stochastic distributed protocol for electric vehicle charging with discrete charging rate," in *IEEE Power and Energy Society General Meeting*, 2012, pp. 1–8.

[28] P. J. Van Laarhoven and E. H. Aarts, *Simulated Annealing*. Springer, 1987.

[29] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-wesley Reading Menlo Park, 1989.

[30] S. H. Strogatz, *Nonlinear Dynamics and Chaos*. Perseus Publishing, 2006.

[31] Southern California Edison, "Average residential load profile of southern california edison," online at http://www.sce.com/005_regul_info/eca/DOMSM11.DLP, 2011.

APPENDIX A
PROOF OF THEOREM 1

We apply the Lasalle Theorem to prove Theorem 1. Note that $x^{(k)}$ depends continuously on $x^{(k-1)}$, the objective function

$$L(x) = \left\| b + \sum_{i=1}^{n} x_i \right\|^2$$

is continuous in x, and the set $\prod_{i=1}^{n} \mathbb{X}_i$ where $\{x^{(k)} : k \geq 1\}$ take values in is compact. To prove Theorem 1, i.e., $x^{(k)} \to \mathbb{O}^*$ as $k \to \infty$, it suffices to prove (i) $L[x^{(k)}] \leq L[x^{(k-1)}]$ and (ii) $L[x^{(k)}] = L[x^{(k-1)}]$ implies $x^{(k-1)} \in \mathbb{O}^*$ for $k \geq 2$. The proofs of (i) and (ii) are presented below.

Fix an arbitrary $k \geq 2$, and let $\Delta x_i := x_i^{(k)} - x_i^{(k-1)}$ denote the update in x_i in iteration k of Algorithm 1 for $i \in \mathcal{N}$. Since $x_i^{(k)}$ minimizes (3) over $\mathbb{X}_i$, it should have a smaller objective value than $x_i^{(k-1)}$, i.e.,

$$2c_i \left\langle g^{(k)}, x_i^{(k)} \right\rangle + \left\| x_i^{(k)} - x_i^{(k-1)} \right\|^2 \leq 2c_i \left\langle g^{(k)}, x_i^{(k-1)} \right\rangle \qquad \Longrightarrow \qquad 2c_i \left\langle g^{(k)}, \Delta x_i \right\rangle + \|\Delta x_i\|^2 \leq 0.$$

Furthermore, since (3) is strictly convex, the equality is attained if and only if $\Delta x_i = 0$.

Abbreviate the aggregate service profile $d^{(k-1)} = b + \sum_{i=1}^{n} x_i^{(k-1)}$ in iteration $k-1$ by d for convenience, then

$$\begin{aligned}
L[x^{(k)}] - L[x^{(k-1)}] = & \left\| d + \sum_{i=1}^{n} \Delta x_i \right\|^2 - \|d\|^2 \\
= & \left\| \sum_{i=1}^{n} \Delta x_i \right\|^2 + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
= & \sum_{i=1}^{n} \|\Delta x_i\|^2 + \sum_{i \neq j} c_i c_j \left\langle \frac{\Delta x_i}{c_i}, \frac{\Delta x_j}{c_j} \right\rangle + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
\leq & \sum_{i=1}^{n} \|\Delta x_i\|^2 + \frac{1}{2} \sum_{i \neq j} c_i c_j \left(\left\| \frac{\Delta x_i}{c_i} \right\|^2 + \left\| \frac{\Delta x_j}{c_j} \right\|^2 \right) + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
= & \sum_{i=1}^{n} \|\Delta x_i\|^2 + \sum_{i \neq j} c_i c_j \left\| \frac{\Delta x_i}{c_i} \right\|^2 + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
= & \sum_{i=1}^{n} \|\Delta x_i\|^2 + \sum_{i=1}^{n} \|\Delta x_i\|^2 \frac{\sum_{j \neq i} c_j}{c_i} + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
= & \sum_{i=1}^{n} \|\Delta x_i\|^2 \frac{\sum_{j=1}^{n} c_j}{c_i} + 2 \sum_{i=1}^{n} \langle d, \Delta x_i \rangle \\
= & \sum_{i=1}^{n} \frac{\sum_{j=1}^{n} c_j}{c_i} \left(\|\Delta x_i\|^2 + 2 c_i \left\langle g^{(k)}, \Delta x_i \right\rangle \right) \ \leq \ 0,
\end{aligned}$$

i.e., (i) holds.

Moreover, it follows from the inequality above that $L[x^{(k)}] = L[x^{(k-1)}]$ if and only if $\Delta x_i = 0$ for $i \in \mathcal{N}$. Therefore if $L[x^{(k)}] = L[x^{(k-1)}]$, then $x_i^{(k-1)} = x_i^{(k)}$ minimizes (3) and it follows from the first order optimality condition that

$$2c_i \left\langle g^{(k)},\ x_i - x_i^{(k-1)} \right\rangle \geq 0,\ i \in \mathcal{N},\ x_i \in \mathbb{X}_i \qquad \Longrightarrow \qquad 2 \left\langle d,\ x_i - x_i^{(k-1)} \right\rangle \geq 0,\ i \in \mathcal{N},\ x_i \in \mathbb{X}_i.$$

The latter is the first order optimality condition for $x^{(k-1)}$ to minimize the objective function $L(x)$. Hence, $L[x^{(k)}] = L[x^{(k-1)}]$ implies $x^{(k-1)} \in \mathbb{O}^*$, i.e., (ii) holds.

We have proved that (i) and (ii) hold for an arbitrary $k \geq 2$, and this completes the proof of Theorem 1.

APPENDIX B
PROOF OF THEOREM 2

We prove (4) by mathematical induction on k.

i) When $k = 0$, $x_i^{(k)} = 0$ for $i = 0, \ldots, n$ and $x_{i'}^{(k)} = 0$ for $i = 1, \ldots, n$, therefore (4) holds.

ii) Assume that (4) holds for $k = K$ ($K \geq 0$). We prove that (4) holds for $k = K+1$ as follows. It is straightforward to verify that $g^{(K+1)}$ is identical for two instances, and therefore

$$\begin{aligned}
x_{i'}^{(K+1)} &= \operatorname*{argmin}_{x_{i'} \in \mathbb{X}_{i'}} 2X_{i'} \left\langle g^{(K+1)}, x_{i'} \right\rangle + \left\| x_{i'} - x_{i'}^{(K)} \right\|^2 \\
&= \operatorname*{argmin}_{x_i \in \mathbb{X}_i} 2X_i \left\langle g^{(K+1)}, x_i \right\rangle + \left\| x_i - x_i^{(K)} \right\|^2 \\
&= x_i^{(K+1)}
\end{aligned}$$

for $i = 2, \ldots, n$ and

$$\begin{aligned} x_{1'}^{(K+1)} &= \operatorname*{argmin}_{x_{1'} \in \mathbb{X}_{1'}} 2X_{1'} \left\langle g^{(K+1)}, x_{1'} \right\rangle + \left\| x_{1'} - x_{1'}^{(K)} \right\|^2 \\ &= \operatorname*{argmin}_{x_{1'} \in 2\mathbb{X}_1} 4X_1 \left\langle g^{(K+1)}, x_{1'} \right\rangle + \left\| x_{1'} - 2x_1^{(K)} \right\|^2 \\ &= 2 \operatorname*{argmin}_{2x_1 \in 2\mathbb{X}_1} 4X_1 \left\langle g^{(K+1)}, 2x_1 \right\rangle + \left\| 2x_1 - 2x_1^{(K)} \right\|^2 \\ &= 2 \operatorname*{argmin}_{x_1 \in \mathbb{X}_1} 2X_1 \left\langle g^{(K+1)}, x_1 \right\rangle + \left\| x_1 - x_1^{(K)} \right\|^2 \\ &= 2x_1^{(K+1)}. \end{aligned}$$

Similarly, $x_{1'}^{(K+1)} = 2X_0^{(K+1)}$. Hence, (4) holds for $k = K + 1$.

According to (i) and (ii), (4) holds for $k \geq 0$, which completes the proof of Theorem 2.

APPENDIX C
PROOF OF THEOREM 3

To prove Theorem 3, we need to bridge a probability distribution P_i over $x_i \in \mathbb{X}_i$ with its expectation $\mathbb{E}_{P_i}[x_i]$. In particular, the following lemma is used when declaring a service profile x is stationary for Algorithm 2.

Lemma 1. *Fix an $i \in \mathcal{N}$ and assume that $\mathbb{X}_i$ satisfies A4, i.e., every element in $\mathbb{X}_i$ has the same ℓ_2 norm $\sqrt{Y_i}$. Let P_i be a probability distribution over $\mathbb{X}_i$, then*

$$\mathbb{E}_{P_i}[x_i] \in \mathbb{X}_i \qquad \Longleftrightarrow \qquad P_i = \delta(x_i) \text{ for some } x_i \in \mathbb{X}_i.$$

Lemma 1 implies that if $\mathbb{E}_{P_i^{(k)}}[x_i] = x_i^{(k-1)}$ for all i, then $P_i^{(k)} = \delta[x_i^{(k-1)}]$ for all i and therefore the tentative charging profile $x^{(k-1)}$ is stationary for Algorithm 2.

Proof: The "$\Leftarrow$" direction is straightforward, and the "$\Rightarrow$" direction can be proved as follows. One has

$$Y_i = \left\| \mathbb{E}_{P_i}[x_i] \right\|^2 \leq \mathbb{E}_{P_i}[\|x_i\|^2] = Y_i$$

and the equality is attained if and only if $P_i = \delta(x_i)$ for some $x_i \in \mathbb{X}_i$. This completes the proof of Lemma 1. ■

Proof of Theorem 3: According to Definition 1, a service profile $x^s = (x_1^s, \ldots, x_n^s)$ is stationary for Algorithm 2 if and only if for any $k \geq 1$, $x^{(k-1)} = x^s$ implies $P_i^{(k)} = \delta(x_i^s)$ for $i \in \mathcal{N}$.

($\Rightarrow$) Let x^s be stationary for Algorithm 2. Fix an arbitrary $k \geq 1$ and assume $x^{(k-1)} = x^s$, then $P_i^{(k)} = \delta(x_i^s)$ for $i \in \mathcal{N}$. Since $P_i^{(k)}$ minimizes (6) over $\Theta(\mathbb{X}_i)$, $x_i^s = \mathbb{E}_{P_i^{(k)}}[x_i]$ minimizes

$$2c_i \left\langle \frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j \neq i} c_j}, x_i \right\rangle + \left\| x_i - x_i^{(k-1)} \right\|^2$$

over $x_i \in \operatorname{conv}(\mathbb{X}_i)$ for $i \in \mathcal{N}$. According to the first order optimality condition,

$$\left\langle 2c_i \frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j \neq i} c_j} + 2\left(x_i^s - x_i^{(k-1)}\right),\ x_i - x_i^s \right\rangle \geq 0$$

for all $x_i \in \mathrm{conv}(\mathbb{X}_i)$ and all $i \in \mathcal{N}$. Substitute $x^{(k-1)} = x^s$ and $g^{(k)} = \left(b + \sum_{j=1}^n x_j^{(k-1)}\right) / \sum_{j=1}^n c_j$ to obtain

$$\left\langle b + \sum_{j \neq i} x_j^s,\ x_i - x_i^s \right\rangle \geq 0$$

for all $x_i \in \mathrm{conv}(\mathbb{X}_i)$ and all $i \in \mathcal{N}$. Then

$$\begin{aligned} \left\langle b + \sum_{j=1}^n x_j^s,\ x_i^s \right\rangle &= \left\langle b + \sum_{j \neq i} x_j^s,\ x_i^s \right\rangle + Y_i \\ &\leq \left\langle b + \sum_{j \neq i} x_j^s,\ x_i \right\rangle + Y_i \\ &= \left\langle b + \sum_{j \neq i} x_j^s + x_i,\ x_i \right\rangle \end{aligned}$$

for all $x_i \in \mathbb{X}_i$ and all $i \in \mathcal{N}$, i.e., x^s is a Nash-equilibrium of the game described in Theorem 3.

($\Leftarrow$) Let x^e be a Nash-equilibrium of the game described in Theorem 3, then

$$\left\langle b + \sum_{j=1}^n x_j^e,\ x_i^e \right\rangle \leq \left\langle b + \sum_{j \neq i} x_j^e + x_i,\ x_i \right\rangle$$

for all $x_i \in \mathbb{X}_i$ and all $i \in \mathcal{N}$. Since $\|x_i^e\|^2 = Y_i = \|x_i\|^2$ for all $x_i \in \mathbb{X}_i$, one has

$$\left\langle b + \sum_{j \neq i} x_j^e,\ x_i^e \right\rangle \leq \left\langle b + \sum_{j \neq i} x_j^e,\ x_i \right\rangle \tag{15}$$

for all $x_i \in \mathbb{X}_i$ and all $i \in \mathcal{N}$. It follows that (15) also holds for all $x_i \in \mathrm{conv}(\mathbb{X}_i)$ and all $i \in \mathcal{N}$. Then,

$$\left\langle b + \sum_{j \neq i} x_j^e,\ x_i - x_i^e \right\rangle \geq 0 \quad \Longrightarrow \quad \left\langle 2c_i \frac{b + \sum_{j \neq i} x_j^e}{\sum_{j \neq i} c_j} + 2\left(x_i^e - x_i^e\right),\ x_i - x_i^e \right\rangle \geq 0$$

for all $x_i \in \mathrm{conv}(\mathbb{X}_i)$ and all $i \in \mathcal{N}$. If $x^{(k-1)} = x^e$ for some $k \geq 1$, then

$$\left\langle 2c_i \frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j \neq i} c_j} + 2\left(x_i^e - x_i^{(k-1)}\right),\ x_i - x_i^e \right\rangle \geq 0$$

for all $x_i \in \mathrm{conv}(\mathbb{X}_i)$ and all $i \in \mathcal{N}$. According to the first order optimality condition, x_i^e minimizes

$$L(x_i) := 2c_i \left\langle \frac{g^{(k)} \sum_{j=1}^n c_j - x_i^{(k-1)}}{\sum_{j \neq i} c_j}, x_i \right\rangle + \left\| x_i - x_i^{(k-1)} \right\|^2 \tag{16}$$

over $x_i \in \mathrm{conv}(\mathbb{X}_i)$ for $i \in \mathcal{N}$. Since (16) is continuous in x_i, x_i^e also minimizes (16) over $\overline{\mathrm{conv}(\mathbb{X}_i)}$ for $i \in \mathcal{N}$.

For each $i \in \mathcal{N}$ and each probability distribution P_i over $\mathbb{X}_i$, let

$$L(P_i) := L(\mathbb{E}_{P_i}[x_i])$$

denote the value of (16) evaluated at the expectation $\mathbb{E}_{P_i}[x_i]$ of P_i, then $L(P_i)$ coincides with the objective in (6). On one hand, since $\mathbb{X}_i$ is compact, the expectation $\mathbb{E}_{P_i}[x_i] \in \overline{\mathrm{conv}(\mathbb{X}_i)}$ and therefore

$$L(P_i) \geq \min_{x_i \in \overline{\mathrm{conv}(\mathbb{X}_i)}} L(x_i) = L(x_i^e)$$

for all $P_i \in \Theta(\mathbb{X}_i)$. On the other hand, the distribution $\delta(x_i^e)$ achieves

$$L[\delta(x_i^e)] = L(x_i^e).$$

Hence, the probability distribution $\delta(x_i^e)$ minimizes (6) over $\Theta(\mathbb{X}_i)$.

Moreover, $\delta(x_i^e)$ is the unique minimizer of $L(P_i)$ over $P_i \in \Theta(\mathbb{X}_i)$. Otherwise, there exists P_i' other than $\delta(x_i^e)$ that minimizes $L(\cdot)$ over $\Theta(\mathbb{X}_i)$. It follows that $L(P_i') = L(\mathbb{E}_{P_i'}[x_i]) = L(x_i^e)$, i.e., both x_i^e and $\mathbb{E}_{P_i'}[x_i]$ minimize $L(x_i)$ over $x_i \in \overline{\text{conv}(\mathbb{X}_i)}$. Since $L(x_i)$ is strictly convex in x_i, one must have

$$\mathbb{E}_{P_i'}[x_i] = x_i^e \in \mathbb{X}_i.$$

It follows from Lemma 1 that $P_i' = \delta(x_i)$ for some $x_i \in \mathbb{X}_i$ and therefore $P_i' = \delta(x_i^e)$. This contradicts our assumption that $P_i' \neq \delta(x_i^e)$. Hence, $\delta(x_i^e)$ is the unique minimizer of $L(P_i)$ over $P_i \in \Theta(\mathbb{X}_i)$.

Hence, $P_i^{(k)} = \delta(x_i^e)$ for $i \in \mathcal{N}$ and therefore x^e is stationary.

Combining ($\Rightarrow$) and ($\Leftarrow$) completes the proof of Theorem 3. ■

APPENDIX D
PROOF OF THEOREM 4

The gap $\|b + \sum_{i=1}^n x_i^*\|^2 - \|b + \sum_{i=1}^n x_i^s\|^2$ can be bounded below as follows.

$$\begin{aligned}
\left\|b + \sum_{i=1}^n x_i^*\right\|^2 - \left\|b + \sum_{i=1}^n x_i^s\right\|^2 &= 2\left\langle b + \sum_{i=1}^n x_i^s,\ \sum_{i=1}^n (x_i^* - x_i^s)\right\rangle + \left\|\sum_{i=1}^n x_i^* - \sum_{i=1}^n x_i\right\|^2 \\
&\geq 2\sum_{i=1}^n \left\langle b + \sum_{j=1}^n x_j^s,\ x_i^* - x_i^s\right\rangle \\
&= 2\sum_{i=1}^n \left\langle b + \sum_{j\neq i} x_j^s,\ x_i^* - x_i^s\right\rangle + 2\sum_{i=1}^n \langle x_i^s,\ x_i^* - x_i^s\rangle \\
&\geq 2\sum_{i=1}^n \langle x_i^s, x_i^* - x_i^s\rangle \qquad \text{[by (15)]} \\
&= 2\sum_{i=1}^n \langle x_i^s, x_i^*\rangle - 2\sum_{i=1}^n Y_i.
\end{aligned}$$

Since $|\langle x_i^s, x_i^*\rangle| \leq \|x_i^s\| \cdot \|x_i^*\| = \sqrt{Y_i} \cdot \sqrt{Y_i} = Y_i$, one has $\langle x_i^s, x_i^*\rangle \geq -Y_i$ for $i \in \mathcal{N}$ and therefore

$$\left\|b + \sum_{i=1}^n x_i^*\right\|^2 - \left\|b + \sum_{i=1}^n x_i^s\right\|^2 \geq 2\sum_{i=1}^n \langle x_i^s, x_i^*\rangle - 2\sum_{i=1}^n Y_i \geq -4\sum_{i=1}^n Y_i.$$

Further, if $\mathbb{X}_i$ is composed of nonnegative elements, then $\langle x_i^s, x_i^*\rangle \geq 0$ for $i \in \mathcal{N}$ and therefore

$$\left\|b + \sum_{i=1}^n x_i^*\right\|^2 - \left\|b + \sum_{i=1}^n x_i^s\right\|^2 \geq 2\sum_{i=1}^n \langle x_i^s, x_i^*\rangle - 2\sum_{i=1}^n Y_i \geq -2\sum_{i=1}^n Y_i.$$

This completes the proof of Theorem 4.

APPENDIX E
PROOF OF PROPOSITION 2

Fix an arbitrary $k \geq 2$ and let $\Delta x_i := x_i^{(k)} - x_i^{(k-1)}$ denote the update of x_i in iteration k of Algorithm 2 for $i \in \mathcal{N}$. Note that the process $\{x^{(k)} : k \geq 0\}$ is Markov. Therefore

$$\mathbb{E}\left[f[x^{(k)}] \mid x^{(k-1)}, x^{(k-2)}, \ldots, x^{(0)}\right] = \mathbb{E}\left[f[x^{(k)}] \mid x^{(k-1)}\right]$$

for any function f. Abbreviate $\mathbb{E}\left[\Delta x_i \mid x^{(k-1)}\right]$ by $\mathbb{E}\Delta x_i$ and abbreviate $\mathbb{E}\left[x_i^{(k)} \mid x^{(k-1)}\right]$ by $\mathbb{E}x_i^{(k)}$ for $i \in \mathcal{N}$ without confusion. Also abbreviate the aggregate service profile $d^{(k-1)} = b + \sum_{i=1}^n x_i^{(k-1)}$ in iteration $k-1$ of Algorithm 2 by d and abbreviate $g^{(k)} = d/\sum_{i=1}^n c_i$ by g, then

$$\begin{aligned}\mathbb{E}\left[L_k \mid x^{(k-1)}\right] - L_{k-1} &= \mathbb{E}\left[\left\|d + \sum_{i=1}^n \Delta x_i\right\|^2 \mid x^{(k-1)}\right] - \|d\|^2 \\ &= \mathbb{E}\left[\left\|\sum_{i=1}^n \Delta x_i\right\|^2 \mid x^{(k-1)}\right] + 2\sum_{i=1}^n \langle d, \mathbb{E}\Delta x_i\rangle .\end{aligned}$$

The first term can be simplified as

$$\begin{aligned}\mathbb{E}\left[\left\|\sum_{i=1}^n \Delta x_i\right\|^2 \mid x^{(k-1)}\right] &= \mathbb{E}\left[\sum_{i=1}^n \|\Delta x_i\|^2 + \sum_{i\neq j}\langle \Delta x_i, \Delta x_j\rangle \mid x^{(k-1)}\right] \\ &= \sum_{i=1}^n \mathbb{E}\left[\|\Delta x_i\|^2 \mid x^{(k-1)}\right] + \sum_{i\neq j}\langle \mathbb{E}\Delta x_i, \mathbb{E}\Delta x_j\rangle\end{aligned}$$

where

$$\begin{aligned}\mathbb{E}\left[\|\Delta x_i\|^2 \mid x^{(k-1)}\right] &= \mathbb{E}\left[\left\|x_i^{(k)} - x_i^{(k-1)}\right\|^2 \mid x^{(k-1)}\right] \\ &= \mathbb{E}\left[\left\|x_i^{(k)}\right\|^2 \mid x^{(k-1)}\right] + \mathbb{E}\left[\left\|x_i^{(k-1)}\right\|^2 \mid x^{(k-1)}\right] - 2\mathbb{E}\left[\left\langle x_i^{(k)}, x_i^{(k-1)}\right\rangle \mid x^{(k-1)}\right] \\ &= 2Y_i - 2\left\langle \mathbb{E}x_i^{(k)}, x_i^{(k-1)}\right\rangle \;=\; -2\left\langle \mathbb{E}\Delta x_i, x_i^{(k-1)}\right\rangle\end{aligned}$$

for $i = 1, 2, \ldots, n$ and

$$\begin{aligned}\sum_{i\neq j}\langle \mathbb{E}\Delta x_i, \mathbb{E}\Delta x_j\rangle &= \sum_{i\neq j} c_i c_j \left\langle \frac{\mathbb{E}\Delta x_i}{c_i}, \frac{\mathbb{E}\Delta x_j}{c_j}\right\rangle \\ &\leq \sum_{i\neq j}\frac{c_i c_j}{2}\left(\frac{\|\mathbb{E}\Delta x_i\|^2}{c_i^2} + \frac{\|\mathbb{E}\Delta x_j\|^2}{c_j^2}\right) \\ &= \sum_{i\neq j} c_i c_j \frac{\|\mathbb{E}\Delta x_i\|^2}{c_i^2} \\ &= \sum_{i=1}^n \frac{\sum_{j\neq i} c_j}{c_i}\|\mathbb{E}\Delta x_i\|^2 .\end{aligned}$$

Therefore

$$\begin{aligned}
\mathbb{E}\left[L_k \mid x^{(k-1)}\right] - L_{k-1} &= \sum_{i=1}^{n}\left[-2\left\langle \mathbb{E}\Delta x_i, x_i^{(k-1)}\right\rangle + \frac{\sum_{j\neq i} c_j}{c_i}\left\|\mathbb{E}\Delta x_i\right\|^2 + 2\left\langle d, \mathbb{E}\Delta x_i\right\rangle\right] \\
&= \sum_{i=1}^{n}\left[\frac{\sum_{j\neq i} c_j}{c_i}\left\|\mathbb{E}\Delta x_i\right\|^2 + 2\left\langle d - x_i^{(k-1)}, \mathbb{E}\Delta x_i\right\rangle\right] \\
&= \sum_{i=1}^{n}\frac{\sum_{j\neq i} c_j}{c_i}\left[\left\|\mathbb{E}\Delta x_i\right\|^2 + 2c_i\left\langle \frac{d - x_i^{(k-1)}}{\sum_{j\neq i} c_j},\ \mathbb{E}\Delta x_i\right\rangle\right] \\
&= \sum_{i=1}^{n}\frac{\sum_{j\neq i} c_j}{c_i}\left[\left\|\mathbb{E}\Delta x_i\right\|^2 + 2c_i\left\langle \frac{g\sum_{j=1}^{n} c_j - x_i^{(k-1)}}{\sum_{j\neq i} c_j},\ \mathbb{E}\Delta x_i\right\rangle\right] \\
&\leq 0 \qquad (17)
\end{aligned}$$

The last inequality is due to the fact that $P_i^{(k)}$ minimizes (6) and therefore should have a smaller objective value than $\delta(x_i^{k-1})$. This completes the proof of Proposition 2.

APPENDIX F
PROOF OF THEOREM 5

The key of proving Theorem 5 is the following lemma, which generalizes the Lassale Theorem from deterministic processes to "hidden" lower bounded supermartingales.

Lemma 2. *Let $\{X_k : k \geq 1\}$ be a time-invariant Markov process that takes values on a compact set $\mathbb{X}$, i.e., X_k is a map from the sample space to $\mathbb{X}$ for $k \geq 1$. Let $L : \mathbb{X} \mapsto \mathbb{R}$ be bounded, i.e., there exists $M \in \mathbb{R}$ such that $|L(x)| \leq M$ for $x \in \mathbb{X}$. Assume that $\{L(X_k) : k \geq 1\}$ is a supermartingale with respect to $\{X_k : k \geq 1\}$, i.e.,*

$$\mathbb{E}\left[L(X_k) \mid X_{k-1}\right] \leq L(X_{k-1}), \qquad k \geq 2.$$

Also assume that the function $f : \mathbb{X} \mapsto \mathbb{R}_-$ defined by

$$f(x) := \mathbb{E}\left[L(X_k) \mid X_{k-1} = x\right] - L(x) \leq 0$$

is continuous and let

$$\mathbb{S} := \{x \in \mathbb{X} \mid f(x) = 0\}$$

denote the set of $x \in \mathbb{X}$ where $\mathbb{E}[L(X_k) \mid X_{k-1} = x] = L(x)$. Then

$$\mathrm{dist}(X_k, \mathbb{S}) \xrightarrow{a.s.} 0, \qquad k \to \infty.$$

Lemma 2 implies that if we can find a function L such that the stochastic process $L(X_k)$ decreases in expectation, then the stochastic process X_k almost surely converges to certain set $\mathbb{S}$, under some technical conditions. Note that if the process $\{X_k : k \geq 1\}$ is deterministic, then Lemma 2 reduces to the Lassale Theorem. The stochastic process $\{X_k : k \geq 1\}$ is called a hidden lower bounded supermartingale since its function $L(X_k)$ is a lower bounded supermartingale. The proof of Lemma 2 is similar to the proof of the Martingale Convergence Theorem [25, Chapter 12.3] and presented below.

Proof: The statement can be transformed as follows.

$$\begin{aligned}
\mathbb{P}\left\{\lim_{k\to\infty}\operatorname{dist}(X_k,\mathbb{S})=0\right\}=1 &\iff \mathbb{P}\left\{\bigcap_{n\geq 1}\left\{\overline{\lim_{k\to\infty}}\operatorname{dist}(X_k,\mathbb{S})<\frac{1}{n}\right\}\right\}=1\\
&\iff \lim_{n\to\infty}\mathbb{P}\left\{\overline{\lim_{k\to\infty}}\operatorname{dist}(X_k,\mathbb{S})<\frac{1}{n}\right\}=1\\
&\iff \mathbb{P}\left\{\overline{\lim_{k\to\infty}}\operatorname{dist}(X_k,\mathbb{S})<\frac{1}{n}\right\}=1,\ \forall n\geq 1\\
&\iff \mathbb{P}\left\{\operatorname{dist}(X_k,\mathbb{S})\geq\frac{1}{n}\text{ infinitely often}\right\}=0,\ \forall n\geq 1\\
&\iff \mathbb{P}\left\{\#\left\{k\mid\operatorname{dist}(X_k,\mathbb{S})\geq\frac{1}{n}\right\}=\infty\right\}=0,\ \forall n\geq 1\\
&\impliedby \mathbb{E}\left[\#\left\{k\mid\operatorname{dist}(X_k,\mathbb{S})\geq\frac{1}{n}\right\}\right]<\infty,\ \forall n\geq 1.
\end{aligned}$$

In the rest of the proof, we show that $\mathbb{E}\left[\#\left\{k\mid\operatorname{dist}(X_k,\mathbb{S})\geq 1/n\right\}\right]<\infty$ for $n\geq 1$.

Fix an arbitrary $n\geq 1$ and let $T:=\#\left\{k\mid\operatorname{dist}(X_k,\mathbb{S})\geq 1/n\right\}$ denote the number of iterations that X_k is at least $1/n$ distance from $\mathbb{S}$, then we want to prove $\mathbb{E}[T]<\infty$. Define an auxiliary process $\{\mathbb{1}_k : k\geq 1\}$ by

$$\mathbb{1}_k:=\begin{cases}1 & \text{if } \operatorname{dist}(X_k,\mathbb{S})\geq 1/n,\\ 0 & \text{otherwise.}\end{cases}$$

Since the set $\mathbb{A}_n:=\{x\in\mathbb{X}\mid\operatorname{dist}(x,\mathbb{S})\geq 1/n\}$ is compact, f is continuous, and $f<0$ on $\mathbb{A}_n$, there exists $\epsilon>0$ such that

$$f(x)\leq-\epsilon,\qquad x\in\mathbb{A}_n.$$

Hence,

$$\begin{aligned}
\mathbb{1}_{k-1}=1 &\quad\Rightarrow\quad \operatorname{dist}(X_{k-1},\mathbb{S})\geq 1/n\\
&\quad\Rightarrow\quad X_{k-1}\in\mathbb{A}_n\\
&\quad\Rightarrow\quad f(X_{k-1})\leq-\epsilon
\end{aligned}$$

for $k\geq 2$ and therefore

$$\begin{aligned}
\mathbb{E}[L(X_K)-L(X_1)] &= \sum_{k=2}^{K}\mathbb{E}[L(X_k)-L(X_{k-1})] = \sum_{k=2}^{K}\mathbb{E}\left[f(X_{k-1})\right]\\
&\leq \sum_{k=2}^{K}\mathbb{E}\left[\mathbb{1}_{k-1}f(X_{k-1})\right] \leq \sum_{k=2}^{K}\mathbb{E}\left[\mathbb{1}_{k-1}\cdot(-\epsilon)\right]\\
&= -\epsilon\sum_{k=1}^{K-1}\mathbb{E}[\mathbb{1}_k]
\end{aligned}$$

for $K\geq 1$. Finally,

$$\mathbb{E}[T]=\mathbb{E}\left[\sum_{k\geq 1}\mathbb{1}_k\right]=\lim_{K\to\infty}\sum_{k=1}^{K-1}\mathbb{E}\left[\mathbb{1}_k\right]\leq\lim_{K\to\infty}\frac{\mathbb{E}[L(X_1)]-\mathbb{E}[L(X_K)]}{\epsilon}\leq\lim_{K\to\infty}\frac{2M}{\epsilon}<\infty.$$

This completes the proof of Lemma 2. ■

Now we apply Lemma 2 to the proof of Theorem 5. In particular, let $X_k = x^{(k)}$ for $k \geq 1$, and let L be the objective function, then $L[x^{(k)}] = L_k$ is a supermartingale. In order to apply Lemma 2, we are left to prove that

$$\mathbb{E}\left[L_k \mid x^{(k-1)}\right] = L_{k-1} \iff x^{(k-1)} \text{ is stationary.}$$

Fix an arbitrary $k \geq 2$, then

$$\begin{aligned}
\mathbb{E}\left[L_k \mid x^{(k-1)}\right] = L_{k-1} &\iff \mathbb{E}[x_i^{(k)} \mid x_i^{(k-1)}] = x_i^{(k-1)} \text{ for } i \in \mathcal{N} && \text{[by (17)]} \\
&\iff \mathbb{E}_{P_i^{(k)}}[x_i] = x_i^{(k-1)} \text{ for } i \in \mathcal{N} \\
&\iff P_i^{(k)} = \delta\left[x_i^{(k-1)}\right] \text{ for } i \in \mathcal{N} && \text{[by Lemma 1]} \\
&\iff x_i^{(k)} = x_i^{(k-1)} \text{ for } i \in \mathcal{N} \\
&\iff x^{(k-1)} \text{ is stationary.}
\end{aligned}$$

This completes the proof of Theorem 5.

APPENDIX G
PROOF OF COROLLARY 2

Since the set $\mathbb{X} := \prod_{i=1}^n \mathbb{X}_i$ has finitely many elements, the minimum distance

$$d_{\min} := \min_{x^1, x^2 \in \mathbb{X}} \|x^1 - x^2\|$$

among different elements in $\mathbb{X}$ is strictly positive. It follows from Theorem 5 that $\operatorname{dist}(x^{(k)}, \mathbb{S}) < d_{\min}$ almost surely as $k \to \infty$.

Pick a random sample path of the process $\{x^{(k)} : k \geq 1\}$. With probability 1, there exists $k_0 \geq 1$ such that $\operatorname{dist}(x^{(k_0)}, \mathbb{S}) < d_{\min}$. It must be that $x^{(k_0)} \in \mathbb{S}$ since otherwise the distance is at least $d_{\min}$. Then, since elements in $\mathbb{S}$ are stationary service profiles, one must have $x^{(k_0)} = x^{(k_0+1)} = \cdots$ and therefore $\lim_{k\to\infty} x^{(k)} = x^{(k_0)} \in \mathbb{S}$ on this sample path.

The sample pathes $\{x^{(k)} : k \geq 1\}$ may converge (with probability 1) to different $x^{(k_0)} \in \mathbb{S}$. Denote the random service profile that the process $\{x^{(k)} : k \geq 1\}$ converges to by x^∞, then

$$x^{(k)} \xrightarrow{\text{a.s.}} x^\infty \text{ as } k \to \infty.$$